

Estimating the Directed Information and Testing for Causality

Ioannis Kontoyiannis, *Fellow, IEEE* *

Maria Skoularidou [†]

April 1, 2016

Abstract

The problem of estimating the directed information rate between two discrete processes $\{X_n\}$ and $\{Y_n\}$ via the plug-in (or maximum-likelihood) estimator is considered. When the joint process $\{(X_n, Y_n)\}$ is a Markov chain of a given memory length, the plug-in estimator is shown to be asymptotically Gaussian and to converge at the optimal rate $O(1/\sqrt{n})$ under appropriate conditions; this is the first estimator that has been shown to achieve this rate. An important connection is drawn between the problem of estimating the directed information rate and that of performing a hypothesis test for the presence of causal influence between the two processes. Under fairly general conditions, the null hypothesis, which corresponds to the absence of causal influence, is equivalent to the requirement that the directed information rate be equal to zero. In that case a finer result is established, showing that the plug-in converges at the faster rate $O(1/n)$ and that it is asymptotically χ^2 -distributed. This is proved by showing that this estimator is equal to (a scalar multiple of) the classical likelihood ratio statistic for the above hypothesis test. Finally it is noted that these results facilitate the design of an actual likelihood ratio test for the presence or absence of causal influence.

Keywords — Entropy, mutual information, directed information, maximum likelihood, plug-in estimator, causality, hypothesis testing, Markov chain, conditional independence, likelihood ratio, χ^2 test

This work was supported by the European Union and Greek National Funds through the Operational Program Education and Lifelong Learning of the National Strategic Reference Framework through the Research Funding Program Thales-Investing in Knowledge Society through the European Social Fund.

Preliminary versions of some of the results in this paper were presented at the Eighth Workshop on Information Theoretic Methods in Science and Engineering, 24-26 June 2015, Copenhagen, Denmark.

*Department of Informatics, Athens University of Economics and Business, Patission 76, Athens 10434, Greece. Email: yiannis@aub.gr.

[†]Department of Informatics, Athens University of Economics and Business, Patission 76, Athens 10434, Greece. Email: m.skoularidou@gmail.com.

1 Introduction

Hypothesis testing and mutual information. One of the most prevalent statistical tools used universally across the sciences today, is the χ^2 test for independence. Suppose we have independent and identically distributed data pairs $(X_1, Y_1), (X_2, Y_2), \dots, (X_n, Y_n)$ and we wish to test whether the X and Y variables are independent or not. Assuming both sets of variables take on finitely many values, we can compare the joint empirical distribution $\hat{P}_{XY,n}(a, b)$ with the product of the empirical marginals $\hat{P}_{X,n}(a)\hat{P}_{Y,n}(b)$; as usual, $\hat{P}_{XY,n}(a, b)$ denotes the proportion of times the pair (a, b) appears in the whole sample, and similarly for the marginals. Pearson's χ^2 test dictates that we compute the (normalized) χ^2 distance between these two distributions,

$$\bar{\chi}_n^2 = n \sum_{a,b} \frac{[\hat{P}_{XY,n}(a, b) - \hat{P}_{X,n}(a)\hat{P}_{Y,n}(b)]^2}{\hat{P}_{X,n}(a)\hat{P}_{Y,n}(b)}. \quad (1)$$

If the data are indeed independent then the distribution of the statistic $\bar{\chi}_n^2$, for large sample sizes n , is approximately χ^2 with $(m-1)(\ell-1)$ degrees of freedom, where m, ℓ are the sizes of the alphabets of X and Y , respectively. Therefore, we can compute the probability of observing a value greater than or equal to $\bar{\chi}_n^2$ under this distribution, and if this probability is appropriately small then we can reject the independence hypothesis.

Another classical test, more closely related to information theoretic ideas, is the likelihood ratio test, based on the statistic,

$$\Delta_n = 2 \log \left(\frac{\prod_{i=1}^n \hat{P}_{XY,n}(X_i, Y_i)}{\prod_{i=1}^n \hat{P}_{X,n}(X_i) \hat{P}_{Y,n}(Y_i)} \right).$$

This has the exact same asymptotic distribution as $\bar{\chi}_n^2$, and an analogous test can be performed. In fact, the $\bar{\chi}_n^2$ statistic can be viewed (and sometimes its use is thus justified) as a quadratic approximation to the nonlinear statistic Δ_n . A more important observation for our purposes is that, after simple algebra, the likelihood ratio test statistic can exactly be expressed as a mutual information,

$$\Delta_n = 2nI(\hat{X}; \hat{Y}) = 2nD(\hat{P}_{XY,n} \| \hat{P}_{X,n} \hat{P}_{Y,n}), \quad (2)$$

where the random variables \hat{X}, \hat{Y} have distribution $\hat{P}_{XY,n}$. Therefore, instead of the χ^2 distance used in (1), the likelihood ratio test statistic (2) examines the (normalized) relative entropy distance between $\hat{P}_{XY,n}$ and $\hat{P}_{X,n}\hat{P}_{Y,n}$. And yet another way to interpret Δ_n is as the “plug-in” estimate of the mutual information $I(X_1; Y_1)$ of the data, using their empirical distribution.

The asymptotic distribution of Δ_n has been re-derived several times historically. In its general form it goes back to the classical result of Wilks [42], see also the texts [23, 35]; in more recent years it has also reappeared in an information-theoretic context, see, e.g., [15, 9, 1].

Estimating directed information and causality testing. This work examines the problem of estimating a different information-theoretic functional: If $\mathbf{X} = \{X_n\}$ and $\mathbf{Y} = \{Y_n\}$ are two finite-valued random process, then the *directed information* $I(X_1^n \rightarrow Y_1^n)$ between $X_1^n = (X_1, X_2, \dots, X_n)$ and $Y_1^n = (Y_1, Y_2, \dots, Y_n)$ is defined as,

$$I(X_1^n \rightarrow Y_1^n) = H(Y_1^n) - \sum_{i=1}^n H(Y_i | Y_1^{i-1}, X_1^i),$$

and the *directed information rate* between \mathbf{X} and \mathbf{Y} is,

$$I(\mathbf{X} \rightarrow \mathbf{Y}) = \lim_{n \rightarrow \infty} \frac{1}{n} I(X_1^n \rightarrow Y_1^n), \quad (3)$$

whenever the limit exists; precise definitions are given in Sections 2 and 3. Directed information was introduced by Massey [25], building on earlier work by Marko [24], in order to provide capacity characterizations for channels with causal feedback. Subsequent work in this direction includes [21, 38, 19, 30, 36, 8], and the dual problem of lossy data compression with feedforward is treated, e.g., in [40]. Additional areas where directed information plays an important role include, among others, distributed and causal data compression and hypothesis testing [14, 29], network communications and control [22, 37], dynamically switching networks [17], sensor networks [33, 45], and causal estimation [41]; see also the references in the above works.

Here we consider the problem of estimating the directed information rate, by tracing the path described above in connection with the mutual information in the reverse direction. Assuming that the joint process $\{(X_n, Y_n)\}$ is a Markov chain of memory length $k \geq 1$, we begin by recalling that under fairly general conditions the limit (3) can be expressed in more manageable form. For example, in Proposition 3.1 we note that when $\{Y_n\}$ itself is also a Markov chain of order no greater than k , the directed information rate is equal to the conditional mutual information $I(\bar{Y}_0; \bar{X}_{-k}^0 | \bar{Y}_{-k}^{-1})$, where $\{(\bar{X}_n, \bar{Y}_n)\}$ denotes the stationary version of the chain.

Main results. In Section 3.2 we define what is probably the simplest estimator of $I(\mathbf{X} \rightarrow \mathbf{Y})$: Interpreting the conditional mutual information $I(\bar{Y}_0; \bar{X}_{-k}^0 | \bar{Y}_{-k}^{-1})$ as a functional of the $(k+1)$ -dimensional distribution of $(\bar{X}_{-k}^0, \bar{Y}_{-k}^0)$, we define the plug-in estimator $\hat{I}_n^{(k)}(\mathbf{X} \rightarrow \mathbf{Y})$ as the same functional of the corresponding empirical distribution. If the original chain is ergodic, then it is easy to see that this estimator is consistent with probability one. Our main results, stated in Theorems 3.2 and 3.3, give finer information for this convergence; cf. (4) and (5) below.

On the one hand, if $I(\mathbf{X} \rightarrow \mathbf{Y}) > 0$ we show that, under appropriate conditions, the plug-in estimator is approximately Gaussian for large n ,

$$\hat{I}_n^{(k)}(\mathbf{X} \rightarrow \mathbf{Y}) \approx N\left(I(\mathbf{X} \rightarrow \mathbf{Y}), \frac{\sigma^2}{n}\right), \quad (4)$$

where the variance σ^2 is identified in Theorem 3.2. Therefore, the plug-in estimator converges at a rate $O(1/\sqrt{n})$ in probability. In fact, Corollary 3.4 shows that the same rate holds in L^1 , and in view of [18, Proposition 3] this implies that the plug-in estimator is optimal in that it converges at the fastest possible rate. Moreover, in Corollary 3.4 we also establish the almost-sure convergence rate of the plug-in estimator, under fairly general conditions.

On the other hand we note that $I(\mathbf{X} \rightarrow \mathbf{Y}) = 0$ if and only if a certain conditional independence property holds, which can be interpreted as the absence of causal influence from \mathbf{X} to \mathbf{Y} : Roughly speaking $I(\mathbf{X} \rightarrow \mathbf{Y})$ is zero if and only if each Y_i , given the past values of the \mathbf{Y} process, is conditionally independent of the values of the \mathbf{X} process up to time i . In fact, one of the main contributions of this work is the identification of two related problems: (1.) understanding the asymptotics of the plug-in estimator $\hat{I}_n^{(k)}(\mathbf{X} \rightarrow \mathbf{Y})$; and (2.) analyzing the likelihood ratio test for the above causality hypothesis.

Intuitively, determining whether the estimated directed information $\hat{I}_n^{(k)}(\mathbf{X} \rightarrow \mathbf{Y})$ is significantly close to zero or not, is related to testing the hypothesis that the above conditional independence relationship holds. Formally, as we show in Proposition 3.5, the (normalized)

likelihood ratio statistic for this test is *exactly* equal to the plug-in estimator $\hat{I}_n^{(k)}(\mathbf{X} \rightarrow \mathbf{Y})$. This connection is described in detail in Section 3.3. Apart from being intellectually satisfying, it also allows us to derive the exact asymptotic distribution of $\hat{I}_n^{(k)}(\mathbf{X} \rightarrow \mathbf{Y})$.

Indeed, we show that, under appropriate conditions, if the directed information rate is zero, then a finer result than (4) can be established for the plug-in estimator; for large n we have,

$$\hat{I}_n^{(k)}(\mathbf{X} \rightarrow \mathbf{Y}) \approx \frac{1}{n} \chi^2(m, \ell, k), \quad (5)$$

where $\chi^2(m, \ell, k)$ is an appropriate χ^2 distribution that only depends on the sizes m, ℓ of the alphabets of X and Y and on the memory size k . In other words, under the null hypothesis (which asserts the presence of the conditional independence being tested), the plug-in estimator converges at a faster rate $O(1/n)$, and the distribution to which it converges only depends on the three known parameters m, ℓ and k . Therefore, a likelihood ratio test for this type of causal influence, can again be performed as before.

In Section 2 we consider the simpler problem of estimating the mutual information rate $\lim_n I(X_n; X_0^{n-1})$ of a Markov chain $\{X_n\}$. The results presented there can be viewed as preliminary versions of the analogous results for the directed information given in Section 3. For clarity of exposition, all proofs are collected in the Appendix.

Earlier work. The connection between the problem of identifying causal relationships and that of testing for conditional independence has a long history. Perhaps the most prominent example is the Granger causality test [11], which uses an autoregressive model (later extended in several directions, most notably to generalized linear models), within which the conditional independence hypothesis described above is tested. The connection of this test with directed information has previously been explored in several directions; see [2] for a comprehensive review.

Several different approaches to the problem of directed information estimation have appeared in the literature in recent years. Rao *et al.* [34] use Miller’s [27] differential entropy estimators in order to estimate (the continuous analog of) directed information, in order to identify causal influence in networks of genes. In the context of neuroscience, Quinn *et al.* [31] use parametric estimation based on generalized linear models to estimate directed information, in order to detect the presence of direct or indirect influence in neuronal networks. And in the subsequent work [32], a near-optimal rate of convergence $O(n^{-1/2+\epsilon})$ is established for the plug-in estimator.

In terms of the present development, the most interesting work is the recent paper by Jiao *et al.* [18]. There, several new estimators for the directed information rate are introduced and they are shown to be consistent under very general conditions. For some of these estimators, particularly those based on the context tree weighting algorithm [43], detailed convergence bounds are also obtained. It is worth noting that our convergence results are obtained under conditions essentially identical to (though slightly weaker than) those required for the bounds in [18]. But using the plug-in also facilitates the connection with hypothesis testing developed here, and makes it possible to obtain, instead of convergence bounds, accurate and sometimes exact asymptotics as described above.

Finally, in a broader context, we note the L^1 and L^2 convergence results presented for mutual information in the very recent work [16]. Although the problems treated there are mostly in the case of independent observations, they provide a general minimax framework for examining the asymptotic optimality of different estimators.

Different approaches. The problem of estimating directed information via the plug-in can and has arisen as a natural question within several different areas. Motivated by applications in econometrics, one of its earliest appearances is in [10], where the directed information functional is defined as a *Kullback causality measure* and is derived as the limiting form of a likelihood ratio statistic used in a temporal causality hypothesis test that is closely related to the test we describe in Section 3.3. For the special case of first-order Markov chains (and under some additional assumptions), the order- $(1/n)$ convergence rate of this statistic to the χ^2 distribution is discussed in [10].

In the physics literature directed information has also appeared extensively under the name *transfer entropy*, and the performance of the corresponding plug-in estimator is examined in the recent works [4, 3]. There, several asymptotic results as well as non-asymptotic bounds are described, though the main emphasis appears to be less on providing rigorous proofs and more on exploring the possible qualitative asymptotic properties of the plug-in.

Arguably the most effective, unifying approach to the task of understanding the behavior of the plug-in estimator for directed information comes from taking the point of view of the asymptotic analysis of maximum likelihood estimates in theoretical statistics. This is indeed the point of view adopted in this work, and for the proofs of our main results we rely on two relevant sets of tools: One is the classical central limit theorem and the law of the iterated logarithm for Markov chains [7], and the other is the general asymptotic theory of statistical inference for Markovian observations [5].

In this vein it should be pointed out (as one of the anonymous reviewers generously outlined at length in their report) that the modern development of classical asymptotics in statistical decision theory provides a powerful technical approach to the task at hand, which can lead to what are apparently the strongest possible results, and accompanied by the sharpest intuition. That development, pioneered by Le Cam, Hajek and others since the 1980s, cf. [6, 39], is built around the “local asymptotic normality,” or LAN, condition. This requires that the log-likelihood ratio between models with nearby parameters, when evaluated at observations produced by a fixed, true distribution corresponding to one of these two parameters, can be appropriately approximated asymptotically by a multivariate Gaussian. For the class of Markov models considered in the present context, this is not hard to verify, as was done, e.g., in [28]. Then, given the LAN condition, the asymptotic behavior of the plug-in, as well its strong optimality properties, can be established by applications of what are, by now, standard results in this area; see the local asymptotic minimax theorem and the convolution theorem in [39, Sections 8.7, 8.9].

2 Mutual information

2.1 Preliminaries

Suppose X is a discrete random variable with values in a finite set A , and with a distribution described by its probability mass function, $P_X(x) = \Pr\{X = x\}$, for $x \in A$. The entropy of X is, $H(X) = H(P_X) = -\sum_{x \in A} P(x) \log P(x)$, where, throughout the paper, ‘log’ denotes the natural logarithm. Viewed as a single random element, the joint entropy of any finite collection of random variables $X_1^n = (X_1, X_2, \dots, X_n)$ is defined analogously; the mutual information between two random variables X and Y is $I(X; Y) = H(X) + H(Y) - H(X, Y)$; the conditional entropy $H(X|Y) = H(X, Y) - H(Y)$; and the conditional mutual information $I(X; Y|Z) = H(X|Z) + H(Y|Z) - H(X, Y|Z)$. As above, we generally write $X_i^j = (X_i, X_{i+1}, \dots, X_j)$, $i \leq j$, for vectors of random variables and similarly $a_i^j = (a_i, a_{i+1}, \dots, a_j) \in A^{j-i+1}$, $i \leq j$, for strings of individual symbols from a finite set A .

The joint distribution of an arbitrary number of discrete random variables is described by their joint probability mass function. For example, the joint distribution of (X, Y, Z) is denoted, $P_{XYZ}(x, y, z) = \Pr\{X = x, Y = y, Z = z\}$. We write the induced marginal distributions in the obvious way, e.g., $P_{XY}(x, y) = \Pr\{X = x, Y = y\}$ and $P_Z(z) = \Pr\{Z = z\}$, and the induced conditionals are similarly denoted, e.g., $P_{XY|Z}(x, y|z) = \Pr\{X = x, Y = y|Z = z\}$.

2.2 The plug-in estimator of mutual information

Suppose $\mathbf{X} = \{X_n\} = \{X_n ; n \geq 0\}$ is a homogeneous, first-order Markov chain on a finite alphabet A , with an arbitrary initial distribution for X_0 and with transition matrix $Q = (Q(a'|a) ; a, a' \in A)$. Although it is not necessary for much of what follows, in order to avoid uninteresting technicalities we assume that $Q(a'|a) > 0$ for all $a, a' \in A$; see the remarks following the statement of Theorem 2.1 regarding how this assumption can be relaxed. Then $\{X_n\}$ has a unique stationary distribution π supported on all of A .

Let $\{\bar{X}_n\}$ denote the stationary version of $\{X_n\}$, having the same transition matrix Q but with initial distribution $\bar{X}_0 \sim \pi$. We wish to estimate the mutual information rate of the Markov chain $\{X_n\}$,

$$\lim_{n \rightarrow \infty} I(X_n; X_0^{n-1}) = \lim_{n \rightarrow \infty} I(X_n; X_{n-1}),$$

which is easily seen to be equal to the mutual information of the chain in equilibrium, namely,

$$I_\pi(X_0; X_1) = I(\bar{X}_0; \bar{X}_1) = I(P_{\bar{X}_0 \bar{X}_1}) = H(\bar{X}_0) + H(\bar{X}_1) - H(\bar{X}_0, \bar{X}_1).$$

Let $\hat{P}_{X_0 X_1, n}$ denote the bivariate empirical distribution obtained from the sample X_0^n ,

$$\hat{P}_{X_0 X_1, n}(a, a') = \frac{1}{n} \sum_{i=1}^n \mathbb{I}_{\{X_{i-1}=a, X_i=a'\}}, \quad a, a' \in A,$$

where \mathbb{I}_E denotes the indicator function of an event E , which equals 1 when E occurs and 0 otherwise. Then the *plug-in estimator* $\hat{I}_n(P_{\bar{X}_0 \bar{X}_1})$ for $I_\pi(X_0; X_1) = I(P_{\bar{X}_0 \bar{X}_1})$ is defined as,

$$\hat{I}_n(P_{\bar{X}_0 \bar{X}_1}) = I(\hat{P}_{X_0 X_1, n}) = H(\hat{P}_{X_0, n}) + H(\hat{P}_{X_1, n}) - H(\hat{P}_{X_0 X_1, n}),$$

where we use the same convention for the notation of marginals, e.g., $\hat{P}_{X_0, n}(a)$, and conditionals, e.g., $\hat{P}_{X_1|X_0, n}(a'|a)$ with those described in Section 2.1.

Recall that, since all the transition probabilities of $\{X_n\}$ are nonzero, the bivariate chain $\{Z_n = (X_{n-1}, X_n) ; n \geq 1\}$ on $A \times A$ is ergodic, so the ergodic theorem [7] implies that, as $n \rightarrow \infty$, $\hat{P}_{X_0, n} \rightarrow \pi$, $\hat{P}_{X_1, n} \rightarrow \pi$, and $\hat{P}_{X_0 X_1, n} \rightarrow P_{\bar{X}_0 \bar{X}_1}$, almost surely (a.s.). Therefore, $\hat{I}_n(P_{\bar{X}_0 \bar{X}_1})$ converges to the desired value, $I(\bar{X}_0; \bar{X}_1)$, a.s., as $n \rightarrow \infty$. The following result describes its finer asymptotic behavior; its proof is given in Appendix A.1.

Theorem 2.1 *Let $\mathbf{X} = \{X_n\}$ be a Markov chain with an all positive transition matrix $Q = (Q(a'|a))$ on the finite alphabet A , and with an arbitrary initial distribution.*

(i) *If the random variables $\{X_n\}$ are not independent, equivalently, if $I_\pi(X_0; X_1) > 0$, then,*

$$\sqrt{n}[\hat{I}_n(P_{\bar{X}_0 \bar{X}_1}) - I_\pi(X_0; X_1)] \xrightarrow{\mathcal{D}} N(0, \sigma^2), \quad \text{as } n \rightarrow \infty,$$

where $\xrightarrow{\mathcal{D}}$ denotes convergence in distribution, $N(0, \sigma^2)$ is the zero-mean normal distribution with variance σ^2 , and with σ^2 given by the following limit, which exists and is finite:

$$\sigma^2 = \lim_{n \rightarrow \infty} \frac{1}{n} \text{Var} \left[\log \left(\prod_{i=1}^n \frac{Q(\bar{X}_i | \bar{X}_{i-1})}{\pi(\bar{X}_i)} \right) \right]. \quad (6)$$

(ii) *If the random variables $\{X_n\}$ are independent, equivalently, if $I_\pi(X_0; X_1) = 0$, then,*

$$2n\hat{I}_n(P_{\bar{X}_0 \bar{X}_1}) \xrightarrow{\mathcal{D}} \chi^2((m-1)^2), \quad \text{as } n \rightarrow \infty,$$

where $\chi^2(s)$ denotes the χ^2 distribution with s degrees of freedom, and $m = |A|$ is the size of the alphabet.

Remarks.

1. As will become evident from the proof, the restriction that all the transition probabilities $Q(a'|a)$ of the chain $\{X_n\}$ are positive, is unnecessary. Indeed, for the result of part (i) it can be entirely removed, and replaced with the minimal assumption that $\{X_n\}$ is irreducible and aperiodic. Similarly, for part (ii) the positivity assumption can be significantly relaxed. For example, of Theorem 5.2 of [5] gives weaker conditions under which the same conclusions can be obtained, e.g., if we restrict attention to a class of ergodic chains whose transition matrices are allowed to contain zero probabilities, but the zeros always occur at the same state transitions.
2. One of the main messages of Theorem 2.1 is the clear dichotomy between independence and dependence: If the random variables $\{X_n\}$ are independent, then $I_\pi(X_0; X_1) = 0$ and the plug-in estimator $\hat{I}_n(P_{\bar{X}_0 \bar{X}_1})$ converges at a rate $O(1/n)$. On the other hand, if the $\{X_n\}$ are not independent, then $I_\pi(X_0; X_1)$ is strictly positive and the plug-in estimator $\hat{I}_n(P_{\bar{X}_0 \bar{X}_1})$ converges at the slower rate $O(1/\sqrt{n})$.

There is a minor caveat in the above syllogism, in that it is only valid as long as σ^2 is strictly positive; when $\sigma^2 = 0$, then even if the $\{X_n\}$ are not independent, the plug-in estimator $\hat{I}_n(P_{\bar{X}_0 \bar{X}_1})$ converges at a rate faster than $O(1/\sqrt{n})$. But it is easy to see, intuitively, that σ^2 is typically nonzero when $\hat{I}_n(P_{\bar{X}_0 \bar{X}_1})$ is positive. In the special case of chains with a uniform stationary distribution, this is illustrated by Proposition 2.2 below, proved in Appendix A.2.

3. As a consequence of the proof of Theorem 2.1, it is fairly simple to determine the exact a.s. rate of convergence of the plug-in estimator under very general conditions. This is stated in Corollary 2.3 below.
4. For the proof of the second part of the theorem we will exploit a connection between the problem of estimating the mutual information $I_\pi(X_0; X_1)$ and a classical hypothesis test for independence, as outlined in the following section.

Proposition 2.2 *Suppose the stationary distribution π of the chain $\{X_n\}$ is uniform on A or, equivalently, that the transition matrix $(Q(a'|a))$ is doubly stochastic. Then the variance σ^2 defined in (6) is zero if and only if the $\{X_n\}$ are i.i.d. and each X_n is uniformly distributed on A .*

The final result of this section gives the exact pointwise rate of convergence for the plug-in estimator; it is established in Appendix A.3.

Corollary 2.3 *Let $\mathbf{X} = \{X_n\}$ be an irreducible and aperiodic Markov chain on the alphabet A , with an arbitrary initial distribution. Then, as $n \rightarrow \infty$, the plug-in estimator satisfies,*

$$\hat{I}_n(P_{\bar{X}_0 \bar{X}_1}) = I_\pi(X_0; X_1) + O\left(\sqrt{\frac{\log \log n}{n}}\right) \quad a.s.$$

2.3 A hypothesis test for independence

Suppose we wish to test the null hypothesis that the random variables $\{X_n\}$ are independent, within the larger hypothesis that $\{X_n\}$ is a Markov chain with all positive transitions. Take, without loss of generality, the alphabet to be $A = \{1, 2, \dots, m\}$, where $m = |A|$. Then, we can parametrize all possible transition matrices $Q = Q_\theta$ with all-positive transition probabilities, by an $m(m-1)$ -dimensional vector θ restricted an open set $\Theta \subseteq \mathbb{R}^{m(m-1)}$. Similarly, the null hypothesis is specified by a lower-dimensional open set $\Phi \subseteq \mathbb{R}^{m(m-1)}$, which is naturally embedded within Θ via a map $h : \Phi \rightarrow \Theta$. The details of the parametrization and the embedding are given in the proof of Theorem 2.1 in Appendix A.1, but, informally, Φ indexes those transition matrices $Q_{h(\phi)}$ that consist of m identical rows, exactly corresponding to those Markov chains that consist of independent random variables $\{X_n\}$.

In order to test the (composite) null hypothesis Φ within the general model Θ , following classical statistical theory we employ a maximum-likelihood ratio test. Specifically, if we define the log-likelihood $L_n(X_0^n; \theta)$ of the sample X_0^n under the distribution corresponding to θ as,

$$L_n(X_0^n; \theta) = \log [\Pr_\theta(X_1^n | X_0)] = \log \left(\prod_{i=1}^n Q_\theta(X_i | X_{i-1}) \right), \quad (7)$$

then the likelihood ratio test statistic is simply the difference,

$$\Delta_n = 2 \left\{ \max_{\theta \in \Theta} L_n(X_0^n; \theta) - \max_{\phi \in \Phi} L_n(X_0^n; h(\phi)) \right\}. \quad (8)$$

In terms of hypothesis testing, there are two important observations to be made here. The first, is that this statistic is *exactly* equal to $2n$ times the plug-in estimator $\hat{I}_n(P_{\bar{X}_0 \bar{X}_1})$; the computation showing this is performed in Appendix A.4.

Proposition 2.4 *Under the assumptions of Theorem 2.1 and in the notation of this section:*

$$\Delta_n = 2n\hat{I}_n(P_{\bar{X}_0\bar{X}_1}).$$

The second important thing to note is that, under the null hypothesis, that is, assuming that the random variables $\{X_n\}$ are independent, part (ii) of Theorem 2.1 tells us that the distribution of $\Delta_n = 2n\hat{I}_n(P_{\bar{X}_0\bar{X}_1})$ is approximately $\chi^2((m-1)^2)$, which does *not* depend on the distribution of the data, except only through the alphabet size m . Therefore, we can decide whether or not the data X_0^n offer strong enough evidence to reject the null hypothesis by examining the value of $\Delta_n = 2n\hat{I}_n(P_{\bar{X}_0\bar{X}_1})$ and then computing a p -value based on this distribution.

Conversely, as we shall see in the proof of (ii) of Theorem 2.1, the asymptotic properties of the estimator $\hat{I}_n(P_{\bar{X}_0\bar{X}_1})$ can be deduced from general results about the likelihood ratio Δ_n .

3 Directed information

3.1 The directed information rate of Markov chains

Let $\mathbf{X} = \{X_n\}$ and $\mathbf{Y} = \{Y_n\}$ be two arbitrary processes with values in the finite alphabets A and B , respectively. Recall that the *directed information* between X_1^n and Y_1^n is,

$$I(X_1^n \rightarrow Y_1^n) = H(Y_1^n) - \sum_{i=1}^n H(Y_i | Y_1^{i-1}, X_1^i),$$

and that it is zero exactly when each Y_i is conditionally independent of X_1^i , given its past Y_1^{i-1} . The natural interpretation of this equivalence is to say that the directed information is zero if and only if \mathbf{X} has no *causal* influence on \mathbf{Y} . We are interested in the problem of estimating the *directed information rate* between \mathbf{X} and \mathbf{Y} , defined as the limit, $I(\mathbf{X} \rightarrow \mathbf{Y}) = \lim_{n \rightarrow \infty} (1/n)I(X_1^n \rightarrow Y_1^n)$, whenever it exists.

If the pairs $\{(X_n, Y_n)\}$ are independent and identically distributed, then it is easy to see that $I(\mathbf{X} \rightarrow \mathbf{Y})$ simplifies to $I(X_1; Y_1)$, and the problem of estimating it reduces to that discussed in the Introduction. Of course, in this case there is nothing to discover regarding causal dependence.

From now on we assume that the pair process $\{(X_n, Y_n) ; n \geq -k + 1\}$ is an ergodic (namely, irreducible and aperiodic) Markov chain on the product alphabet $A \times B$, of memory length $k \geq 1$, and with an arbitrary initial distribution for (X_{-k+1}^0, Y_{-k+1}^0) . We write $\{(\bar{X}_n, \bar{Y}_n)\}$ for the stationary version of $\{(X_n, Y_n)\}$ with (X_{-k+1}^0, Y_{-k+1}^0) distributed according to the unique invariant measure of the bivariate chain, and recall that the distribution of $\{(\bar{X}_n, \bar{Y}_n)\}$ can be extended so that it is defined for all $n = \dots, -1, 0, 1, \dots$.

In this case, the following proposition shows that, under appropriate conditions, the directed information rate can be expressed in simpler form.

Proposition 3.1 *Suppose $\{(X_n, Y_n)\}$ is an irreducible and aperiodic Markov chain of order no larger than k , with an arbitrary initial distribution. Then:*

- (i) *The entropy rate $H(\mathbf{Y})$ of the univariate process $\mathbf{Y} = \{Y_n\}$ exists and,*

$$H(\mathbf{Y}) = \lim_{n \rightarrow \infty} \frac{1}{n} H(Y_1^n) = \lim_{n \rightarrow \infty} \frac{1}{n} H(\bar{Y}_1^n).$$

(ii) The directed information rate $I(\mathbf{X} \rightarrow \mathbf{Y})$ exists and it equals,

$$I(\mathbf{X} \rightarrow \mathbf{Y}) = \lim_{n \rightarrow \infty} \frac{1}{n} I(X_1^n \rightarrow Y_1^n) = H(\mathbf{Y}) - H(\bar{Y}_0 | \bar{X}_{-k}^0, \bar{Y}_{-k}^{-1}).$$

(iii) If $\mathbf{Y} = \{Y_n\}$ is also a Markov chain of order no larger than k , then $I(\mathbf{X} \rightarrow \mathbf{Y})$ further simplifies to,

$$I(\mathbf{X} \rightarrow \mathbf{Y}) = I(\bar{Y}_0; \bar{X}_{-k}^0 | \bar{Y}_{-k}^{-1}).$$

Remarks.

1. Throughout this section we assume that $\{(X_n, Y_n)\}$ is a Markov chain, not necessarily stationary (i.e., with an arbitrary initial distribution), with memory no larger than some fixed k . For the sake of technical convenience we will also assume that $\{(X_n, Y_n)\}$ has a strictly positive transition matrix Q ,

$$Q(a_k, b_k | a_0^{k-1}, b_0^{k-1}) = \Pr\{X_n = a_k, Y_n = b_k | X_{n-k}^{n-1} = a_0^{k-1}, Y_{n-k}^{n-1} = b_0^{k-1}\} > 0,$$

for all $a_0^k \in A^{k+1}$, $b_0^k \in B^{k+1}$. As discussed in the remarks following Theorem 3.3, this assumption can be significantly relaxed.

2. Like mutual information, the directed information rate $I(\mathbf{X} \rightarrow \mathbf{Y})$ also admits important operational interpretations. For example, in the case of a stationary k th order Markov chain $\{(X_n, Y_n)\}$ such that $\{Y_n\}$ is also a k th order chain, we can use the data processing property of mutual information in the result of part (iii) of the proposition to see that,

$$I(\mathbf{X} \rightarrow \mathbf{Y}) = I(Y_0; X_{-k}^0 | Y_{-k}^{-1}) = I(Y_0; X_{-\infty}^0 | Y_{-\infty}^{-1}).$$

This quantity is zero if and only if each Y_i , given its past $Y_{-\infty}^{i-1}$, is conditionally independent of $X_{-\infty}^i$, confirming our original intuition that the directed information is only zero in the absence of causal influence.

3. In the case of a general stationary chain $\{(X_n, Y_n)\}$, without assuming anything else about the process $\{Y_n\}$, data processing still implies that,

$$I(Y_0; X_{-k}^0 | Y_{-k}^{-1}) = I(Y_0; X_{-\infty}^0 | Y_{-k}^{-1}) \geq I(Y_0; X_{-\infty}^0 | Y_{-\infty}^{-1}).$$

This is zero if and only if Y_0 , given only its k -past Y_{-k}^{-1} , is conditionally independent of $X_{-\infty}^0$. In this case the quantity $I(Y_0; X_{-k}^0 | Y_{-k}^{-1})$ is not enough to entirely characterize the absence of causal influence from \mathbf{X} to \mathbf{Y} , but knowing its value nevertheless offers some evidence for such an influence. In particular, knowing that it is zero (or sufficiently close to zero), would still imply that \mathbf{X} has no (or little) causal influence on \mathbf{Y} .

Therefore, even if \mathbf{Y} is not necessarily Markovian, it is always of interest to estimate $I(\bar{Y}_0; \bar{X}_{-k}^0 | \bar{Y}_{-k}^{-1})$. Indeed, as we explain in detail in Section 3.3, this estimation problem is intimately related to a likelihood-ratio hypothesis test for the presence of causal influence.

3.2 The plug-in estimator of directed information rate

Given a sample (X_{-k+1}^n, Y_{-k+1}^n) from the joint process $\{(X_n, Y_n)\}$, we define the $(k+1)$ -dimensional, bivariate empirical distribution induced on $A^{k+1} \times B^{k+1}$, as,

$$\hat{P}_{X_{-k}^0 Y_{-k}^0, n}(a_0^k, b_0^k) = \frac{1}{n} \sum_{i=1}^n \mathbb{I}_{\{X_{i-k}^i = a_0^k, Y_{i-k}^i = b_0^k\}}, \quad a_0^k \in A^{k+1}, b_0^k \in B^{k+1}. \quad (9)$$

Motivated by the discussion in the remarks following Proposition 3.1, we now define *the plug-in estimator* for the directed information rate $I(\mathbf{X} \rightarrow \mathbf{Y})$ as,

$$\hat{I}_n^{(k)}(\mathbf{X} \rightarrow \mathbf{Y}) = I(\hat{Y}_0; \hat{X}_{-k}^0 | \hat{Y}_{-k}^{-1}), \quad \text{where } (\hat{X}_{-k}^0, \hat{Y}_{-k}^0) \sim \hat{P}_{X_{-k}^0 Y_{-k}^0, n}. \quad (10)$$

Since all the transition probabilities of the bivariate chain $\{(X_n, Y_n)\}$ are nonzero, the $(k+1)$ -dimensional chain $\{Z_n = (X_{n-k}^n, Y_{n-k}^n)\}$ is ergodic, so the ergodic theorem [7] implies that the empirical distributions $\hat{P}_{X_{-k}^0 Y_{-k}^0, n}$ converge a.s., as $n \rightarrow \infty$, to $P_{\bar{X}_{-k}^0 \bar{Y}_{-k}^0}$. And hence, the plug-in estimator $\hat{I}_n^{(k)}(\mathbf{X} \rightarrow \mathbf{Y})$ also converges a.s. to the desired value, $I(\bar{Y}_0; \bar{X}_{-k}^0 | \bar{Y}_{-k}^{-1})$. The following result describes its finer asymptotic behavior.

Theorem 3.2 *Let $\{(X_n, Y_n)\}$ be a Markov chain of memory length $k \geq 1$, with an all positive transition matrix Q on the finite alphabet $A \times B$, and with an arbitrary initial distribution. Assume that the univariate process $\{Y_n\}$ is also a Markov chain with memory length k .*

- (i) *If the random variables $\{X_n\}$ do have a causal influence on the $\{Y_n\}$, equivalently, if $I(\mathbf{X} \rightarrow \mathbf{Y}) > 0$ then,*

$$\sqrt{n} [\hat{I}_n^{(k)}(\mathbf{X} \rightarrow \mathbf{Y}) - I(\mathbf{X} \rightarrow \mathbf{Y})] \xrightarrow{\mathcal{D}} N(0, \sigma^2), \quad \text{as } n \rightarrow \infty,$$

where the variance σ^2 is given by the following limit, which exists and is finite:

$$\sigma^2 = \lim_{n \rightarrow \infty} \frac{1}{n} \text{Var} \left\{ \log \left[\prod_{i=1}^n \left(\frac{P_{\bar{X}_{-k}^0 \bar{Y}_0 | \bar{Y}_{-k}^{-1}}(X_{i-k}^i, Y_i | Y_{i-k}^{i-1})}{P_{\bar{Y}_0 | \bar{Y}_{-k}^{-1}}(Y_i | Y_{i-k}^{i-1}) P_{\bar{X}_{-k}^0 | \bar{Y}_{-k}^{-1}}(X_{i-k}^i | Y_{i-k}^{i-1})} \right) \right] \right\}. \quad (11)$$

- (ii) *If the random variables $\{X_n\}$ do not have a causal influence on the $\{Y_n\}$, equivalently, if $I(\mathbf{X} \rightarrow \mathbf{Y}) = 0$ then,*

$$2n \hat{I}_n^{(k)}(\mathbf{X} \rightarrow \mathbf{Y}) \xrightarrow{\mathcal{D}} \chi^2 \left(\ell^k (m^{k+1} - 1) (\ell - 1) \right), \quad \text{as } n \rightarrow \infty, \quad (12)$$

where $m = |A|$ and $\ell = |B|$ are the sizes of the alphabets A, B , respectively.

Theorem 3.2 is an immediate consequence of the following more general result that does not assume that \mathbf{Y} is a Markov chain, combined with Proposition 3.1. Theorem 3.3 is proved in Appendix A.6.

Theorem 3.3 *If $\{(X_n, Y_n)\}$ is a Markov chain of memory length $k \geq 1$, with an all positive transition matrix Q on the finite alphabet $A \times B$, and with an arbitrary initial distribution, then:*

(i) *If $I(\bar{Y}_0; \bar{X}_{-k}^0 | \bar{Y}_{-k}^{-1})$ is nonzero, then with σ^2 as in (11):*

$$\sqrt{n} \left[\hat{I}_n^{(k)}(\mathbf{X} \rightarrow \mathbf{Y}) - I(\bar{Y}_0; \bar{X}_{-k}^0 | \bar{Y}_{-k}^{-1}) \right] \xrightarrow{\mathcal{D}} N(0, \sigma^2), \quad \text{as } n \rightarrow \infty.$$

(ii) *If, on the other hand, $I(\bar{Y}_0; \bar{X}_{-k}^0 | \bar{Y}_{-k}^{-1}) = 0$, then the plug-in estimator converges to a χ^2 distribution exactly as in (12).*

Remarks.

1. From the proof of Theorem 3.3 it is evident that the restriction of all-positive transition probabilities $Q(a_k, b_k | a_0^{k-1}, b_0^{k-1})$ for the chain $\{(X_n, Y_n)\}$ is unnecessary. The result of part (i) remains valid with this restriction replaced with the minimal assumption that the pair process $\{(X_n, Y_n)\}$ is irreducible and aperiodic. And for part (ii) the positivity assumption can also be significantly relaxed, in accordance with the discussion around Theorem 5.2 of [5], particularly as long as the k -dimensional version of the assumptions in Condition 5.1 is satisfied, as discussed in Remark 1 after Theorem 2.1 earlier.
2. An important consequence of Theorems 3.2 and 3.3 is the clear dichotomy between the presence and absence of causal influence: If the $\{X_n\}$ have no causal influence on the $\{Y_n\}$, then $I(\mathbf{X} \rightarrow \mathbf{Y}) = 0$ and the plug-in estimator converges at a rate $O(1/n)$. On the other hand, if such causal influence does exist, then the directed information rate $I(\mathbf{X} \rightarrow \mathbf{Y})$ is strictly positive and the plug-in estimator converges at the slower rate $O(1/\sqrt{n})$.
3. An examination of the proof of Theorem 3.3 shows that, with some additional effort, it can be refined to provide very accurate results on the a.s. and L^1 rates of convergence of the plug-in estimator, under very general conditions; see Corollary 3.4 below, proved in Appendix A.7. In particular, in view of the converse result in [18, Proposition 3], the asymptotic bound in (14) implies that the L^1 rate at which the plug-in converges is optimal.

Although it is easily established that an analogous result holds for the plug-in estimator of mutual information, the corresponding proof is merely a simplification of the (already fairly straightforward) proof of Corollary 3.4, and therefore it was not included in the previous section.

4. For the proof of the χ^2 convergence part of the theorem we will exploit an interesting connection of this problem with a classical hypothesis test for conditional independence; this is discussed in detail in Section 3.3.

Corollary 3.4 *Let $\{(X_n, Y_n)\}$ be an irreducible and aperiodic Markov chain on $A \times B$, of memory length $k \geq 1$, and with an arbitrary initial distribution. Then, as $n \rightarrow \infty$, the plug-in estimator satisfies, as $n \rightarrow \infty$,*

$$\hat{I}_n^{(k)}(\mathbf{X} \rightarrow \mathbf{Y}) - I(\bar{Y}_0; \bar{X}_{-k}^0 | \bar{Y}_{-k}^{-1}) = O\left(\sqrt{\frac{\log \log n}{n}}\right) \quad \text{a.s.}, \quad (13)$$

$$E \left[|\hat{I}_n^{(k)}(\mathbf{X} \rightarrow \mathbf{Y}) - I(\bar{Y}_0; \bar{X}_{-k}^0 | \bar{Y}_{-k}^{-1})| \right] = O\left(\frac{1}{\sqrt{n}}\right). \quad (14)$$

If \mathbf{Y} is also a Markov chain of memory no larger than k , then, as $n \rightarrow \infty$,

$$\begin{aligned}\hat{I}_n^{(k)}(\mathbf{X} \rightarrow \mathbf{Y}) - I(\mathbf{X} \rightarrow \mathbf{Y}) &= O\left(\sqrt{\frac{\log \log n}{n}}\right) \quad a.s. \\ E\left[|\hat{I}_n^{(k)}(\mathbf{X} \rightarrow \mathbf{Y}) - I(\mathbf{X} \rightarrow \mathbf{Y})|\right] &= O\left(\frac{1}{\sqrt{n}}\right).\end{aligned}\quad (15)$$

In view of [18, Proposition 3], the L^1 convergence rate established in (15) above is optimal.

3.3 A hypothesis test for causal influence

Suppose we wish to test whether or not the samples $\{X_n\}$ have a causal influence on the $\{Y_n\}$. As discussed already, in the present context this translates to testing the null hypothesis that each random variable Y_i is conditionally independent of X_{i-k}^i given Y_{i-k}^{i-1} , within the larger hypothesis that the pair process $\{(X_n, Y_n)\}$ is a k th order Markov chain on $A \times B$ with all positive transitions. We take, without loss of generality, the alphabets of \mathbf{X} and \mathbf{Y} to be $A = \{1, 2, \dots, m\}$ and $B = \{1, 2, \dots, \ell\}$, respectively.

As we describe in detail in the proof of Theorem 3.3 in Appendix A.6, each positive transition matrix $Q = Q_\theta$ is indexed by a parameter vector θ taking values in an $m^k \ell^k (m\ell - 1)$ -dimensional open set Θ . And the null hypothesis corresponding to each random variable Y_i being conditionally independent of X_{i-k}^i given Y_{i-k}^{i-1} , is described by transition matrices Q_θ which can be decomposed as,

$$Q_\theta(a_0, b_0 | a_{-k}^{-1}, b_{-k}^{-1}) = Q_\theta^x(a_0 | a_{-k}^{-1}, b_{-k}^{-1}) Q_\theta^y(b_0 | b_{-k}^{-1}). \quad (16)$$

This is formally described by a lower-dimensional parameter set Φ , which can be embedded in Θ via a map $h : \Phi \rightarrow \Theta$, such that all induced transition matrices $Q_{h(\phi)}$ correspond to Markov chains that satisfy the required conditional independence property (16).

In order to test the null hypothesis Φ within the general model Θ , we employ a likelihood ratio test. Specifically, we define the log-likelihood $L_n(X_{-k+1}^n, Y_{-k+1}^n; \theta)$ of the sample (X_{-k+1}^n, Y_{-k+1}^n) under the distribution corresponding to θ as,

$$L_n(X_{-k+1}^n, Y_{-k+1}^n; \theta) = \log [\Pr_\theta(X_1^n, Y_1^n | X_{-k+1}^0, Y_{-k+1}^0)] = \log \left(\prod_{i=1}^n Q_\theta(X_i, Y_i | X_{i-k}^{i-1}, Y_{i-k}^{i-1}) \right), \quad (17)$$

so that the likelihood ratio test statistic is simply the difference,

$$\Delta_n = 2 \left\{ \max_{\theta \in \Theta} L_n(X_{-k+1}^n, Y_{-k+1}^n; \theta) - \max_{\phi \in \Phi} L_n(X_{-k+1}^n, Y_{-k+1}^n; h(\phi)) \right\}. \quad (18)$$

As in the earlier case discussed in Section 2.3, there are two key observation to be made here. First, this statistic is *exactly* equal to $2n$ times the plug-in estimator. Proposition 3.5 is proved in Appendix A.8.

Proposition 3.5 *Under the assumptions of Theorem 3.3 and in the notation of this section:*

$$\Delta_n = 2n \hat{I}_n^{(k)}(\mathbf{X} \rightarrow \mathbf{Y}).$$

Recall that, under the null hypothesis, part (ii) of Theorem 3.3 tells us that the distribution of Δ_n is approximately χ^2 with $\ell^k(m^{k+1} - 1)(\ell - 1)$ degrees of freedom. The second important thing to note is that this limiting distribution does not depend on the actual distribution of the samples, except through the alphabet sizes m, ℓ and the memory length k . Therefore, following standard statistical methodology, we can decide whether or not the data offer strong enough evidence to reject the null hypothesis by examining the value of Δ_n : If the p -value given by the probability of the tail $[\Delta_n, \infty)$ of the asymptotic χ^2 distribution is below a certain threshold α , then the causality hypothesis can be rejected at the significance level α .

Conversely, as we discuss in the proof of part (ii) of Theorem 3.3, the asymptotic distribution of the plug-in estimator under the null hypothesis, follows from the corresponding general results about the likelihood ratio in [5].

A Appendix

A.1 Proof of Theorem 2.1

For part (i), first, we express $\hat{I}_n(P_{\bar{X}_0\bar{X}_1}) = I(\hat{P}_{X_0X_1,n})$ as,

$$\begin{aligned}
& \sum_{a,a'} \hat{P}_{X_0X_1,n}(a,a') \log \left(\frac{\hat{P}_{X_0X_1,n}(a,a')}{\hat{P}_{X_0,n}(a)\hat{P}_{X_1,n}(a')} \right) \\
&= \sum_{a,a'} \hat{P}_{X_0X_1,n}(a,a') \log \left(\frac{Q(a'|a)}{\pi(a')} \right) + \sum_{a,a'} \hat{P}_{X_0X_1,n}(a,a') \log \left(\frac{\pi(a')}{Q(a'|a)} \frac{\hat{P}_{X_1|X_0,n}(a'|a)}{\hat{P}_{X_1,n}(a')} \right) \\
&= \sum_{a,a'} \left[\left(\frac{1}{n} \sum_{i=1}^n \mathbb{I}_{\{X_{i-1}=a, X_i=a'\}} \right) \log \left(\frac{Q(a'|a)}{\pi(a')} \right) \right] \\
&\quad + \sum_{a,a'} \hat{P}_{X_0X_1,n}(a,a') \log \left(\frac{\pi(a)}{\hat{P}_{X_0,n}(a)} \frac{\pi(a')}{\hat{P}_{X_1,n}(a')} \frac{\hat{P}_{X_0X_1,n}(a,a')}{Q(a'|a)\pi(a)} \right) \\
&= \frac{1}{n} \sum_{i=1}^n \left[\sum_{a,a'} \left(\mathbb{I}_{\{X_{i-1}=a, X_i=a'\}} \log \left(\frac{Q(a'|a)}{\pi(a')} \right) \right) \right] \\
&\quad + D(\hat{P}_{X_0X_1,n} \| P_{\bar{X}_0\bar{X}_1}) - D(\hat{P}_{X_0,n} \| \pi) - D(\hat{P}_{X_1,n} \| \pi) \\
&= \frac{1}{n} \sum_{i=1}^n \log \left(\frac{Q(X_i|X_{i-1})}{\pi(X_i)} \right) + D(\hat{P}_{X_0X_1,n} \| P_{\bar{X}_0\bar{X}_1}) - D(\hat{P}_{X_0,n} \| \pi) - D(\hat{P}_{X_1,n} \| \pi), \quad (19)
\end{aligned}$$

where $D(P\|Q) = \sum_{x \in A} P(x) \log[P(x)/Q(x)]$ denotes the relative entropy between two discrete distributions P and Q on the same alphabet A . As mentioned earlier, the bivariate chain $\{Z_n = (X_{n-1}, X_n) ; n \geq 1\}$ on $A \times A$ is ergodic, therefore it satisfies the ergodic theorem, the central limit theorem, and the law of the iterated logarithm; see, e.g., [7]. In particular, the law of the iterated logarithm implies that, each of the three relative entropies above, when multiplied by \sqrt{n} , converges to zero a.s., as $n \rightarrow \infty$. Indeed, a Taylor expansion shows that $-\sqrt{n}D(\hat{P}_{X_0,n} \| \pi)$ is equal to,

$$\begin{aligned}
& \sqrt{n} \sum_{a \in A} \hat{P}_{X_0,n}(a) \log \left[1 + \left(\frac{\pi(a)}{\hat{P}_{X_0,n}(a)} - 1 \right) \right] \\
&= \sqrt{n} \sum_{a \in A} \hat{P}_{X_0,n}(a) \left[\left(\frac{\pi(a)}{\hat{P}_{X_0,n}(a)} - 1 \right) - \frac{1}{2} \left(\frac{\pi(a)}{\hat{P}_{X_0,n}(a)} - 1 \right)^2 \frac{1}{\xi_n(a)^2} \right] \\
&= -\frac{\sqrt{n}}{2} \sum_{a \in A} \hat{P}_{X_0,n}(a) \left[\left(\frac{\pi(a)}{\hat{P}_{X_0,n}(a)} - 1 \right)^2 \frac{1}{\xi_n(a)^2} \right] \\
&= -\frac{\sqrt{n}}{2} \sum_{a \in A} \left[\frac{1}{\hat{P}_{X_0,n}(a)\xi_n(a)^2} (\hat{P}_{X_0,n}(a) - \pi(a))^2 \right],
\end{aligned}$$

for some (random) $\xi_n(a)$ between 1 and $\pi(a)/\hat{P}_{X_0,n}(a)$. By the ergodic theorem, $(1/\hat{P}_{X_0,n}(a))$ and $(1/\xi_n(a)^2)$ are both bounded a.s., and from the law of the iterated logarithm we have that

$(\hat{P}_{X_0,n}(a) - \pi(a))\sqrt{n/(\log \log n)}$ is also bounded a.s. Therefore, writing the above expression as,

$$- \frac{\log \log n}{2\sqrt{n}} \sum_{a \in A} \left[\frac{1}{\hat{P}_{X_0,n}(a)\xi_n(a)^2} (\hat{P}_{X_0,n}(a) - \pi(a))^2 \left(\frac{n}{\log \log n} \right) \right], \quad (20)$$

we see that each term in the sum is a.s. bounded, and, therefore, the entire expression tends to zero, a.s., as $n \rightarrow \infty$, as required.

The same argument shows that, after multiplication by \sqrt{n} , the other two relative entropies in (19) also tend to zero, a.s., as $n \rightarrow \infty$. Therefore, a.s. as $n \rightarrow \infty$,

$$\sqrt{n}[\hat{I}_n(P_{\bar{X}_0\bar{X}_1}) - I_\pi(X_0; X_1)] = \frac{1}{\sqrt{n}} \sum_{i=1}^n \left[\log \left(\frac{Q(X_i|X_{i-1})}{\pi(X_i)} \right) - I_\pi(X_0; X_1) \right] + o(1), \quad (21)$$

and the result of part (i) follows immediately by an application of the central limit theorem [7, Sec. I.16] to the bivariate chain $\{Z_n\}$. The existence of the limit in the definition of σ^2 is guaranteed by [7, Theorem 3, p. 97], and its finiteness follows easily from the fact that the alphabet is finite and the random variables $\{\log(Q(\bar{X}_i|\bar{X}_{i-1})/\pi(\bar{X}_i))\}$ are uniformly bounded.

For part (ii), first recall the hypothesis testing setup presented in Section 2.3. Each transition matrix $Q = Q_\theta$ on A with positive transitions is indexed by a parameter θ in the following open subset of $\mathbb{R}^{m(m-1)}$,

$$\Theta = \left\{ \theta \in \mathbb{R}^{m(m-1)} : \theta_{ij} > 0 \text{ for all } i, j, \text{ and } \sum_{1 \leq j \leq m-1} \theta_{ij} < 1 \text{ for each } i \right\}.$$

The entries of the transition matrix Q_θ corresponding to a given parameter $\theta \in \Theta$ are,

$$Q_\theta(j|i) = \begin{cases} \theta_{ij}, & \text{if } j \neq m \\ 1 - \sum_{1 \leq j' \leq m-1} \theta_{ij'}, & \text{if } j = m \end{cases}, \quad \text{for all } i, j. \quad (22)$$

Similarly, the null hypothesis is specified by the open set,

$$\Phi = \left\{ \phi \in \mathbb{R}^{m-1} : \phi_j > 0 \text{ for all } j, \text{ and } \sum_{1 \leq j \leq m-1} \phi_j < 1 \right\},$$

which is naturally embedded within Θ via the map $h : \Phi \rightarrow \Theta$, where $\phi \mapsto \theta = h(\phi)$, with $\theta_{ij} = \phi_j$ for all i, j .

Now, recall the result of Proposition 2.4, stating that our quantity of interest, $2n\hat{I}_n(P_{\bar{X}_0\bar{X}_1})$, is equal to twice the log-likelihood ratio Δ_n defined in (8). Then, the claimed convergence of $2n\hat{I}_n(P_{\bar{X}_0\bar{X}_1})$ is exactly the result stated as (the last part of) the conclusion of [5, Theorem 5.2]. For that, we only need to verify the two assumptions of that theorem.

For the first assumption, we note that, in the notation of [5], the size s of the alphabet $s = m$, the dimensionality r of Θ is $r = m(m-1)$, and the dimensionality c of Φ is $c = m-1$, so that the limiting distribution of Δ_n has $r - c = (m-1)^2$ degrees of freedom. For each (i, j) , $1 \leq i \leq m$, $1 \leq j \leq m-1$, the $(i, j)^{th}$ component of h given by $h_{ij}(\phi) = \phi_j$ is certainly three times continuously differentiable with respect to each ϕ_ℓ , $1 \leq \ell \leq m-1$. Moreover, the $m(m-1) \times (m-1)$ matrix $K(\phi)$ with entries,

$$(K(\phi))_{ij,\ell} = \frac{\partial h_{ij}(\phi)}{\partial \phi_\ell} = \frac{\partial \phi_j}{\partial \phi_\ell} = \delta_{j\ell},$$

where $\delta_{j\ell} = \delta_{j,\ell}$ equals 1 if $j = \ell$ and 0 otherwise, has rank $c = m - 1$ throughout Φ . This shows that the first assumption we needed, namely, Condition 3.1 on [5, p. 17], indeed is satisfied.

Similarly, for the second assumption let $D = A \times A$ and write $d = |D| = m^2$. For each (i, j) , $1 \leq i, j \leq m$, the component of the transition matrix Q_θ corresponding to (i, j) , $Q_\theta(j|i)$, is certainly three times continuously differentiable throughout Θ . Moreover, recalling the parametrization (22), consider the $m^2 \times m(m-1)$ matrix $\mathcal{K}(\theta)$ with entries, for $1 \leq i, j, \ell \leq m$, $1 \leq k \leq m-1$,

$$(\mathcal{K}(\theta))_{ij, \ell k} = \frac{\partial Q_\theta(j|i)}{\partial \theta_{\ell k}} = \begin{cases} \delta_{ij, \ell k}, & \text{if } j \neq m; \\ -\delta_{i, \ell}, & \text{if } j = m. \end{cases}$$

Clearly $\mathcal{K}(\theta)$ has rank $r = m(m-1)$ throughout Θ , which shows that the second assumption we needed to check, namely Condition 5.1 on [5, p. 23], is also satisfied, thus completing the proof. \square

A.2 Proof of Proposition 2.2

This result is a consequence of [20, Theorem 3]; see also the discussion in [44]. To see that, observe that, under the assumptions of the proposition, $\pi(\bar{X}_i) = 1/|A|$ for all i and therefore,

$$\sigma^2 = \lim_{n \rightarrow \infty} \frac{1}{n} \text{Var} \left[\log \left(\prod_{i=1}^n Q(\bar{X}_i | \bar{X}_{i-1}) \right) \right],$$

which is equal to the variance σ^2 defined in [20, eq. (3.2)]. The stronger version of Theorem 3 established in the last section of [20] implies that σ^2 is only zero when there is a constant $q > 0$ and a positive vector $(v(a) ; a \in A)$ such that,

$$Q(a'|a) = q \frac{v(a')}{v(a)}, \quad \text{for all } a, a' \in A.$$

But since all $Q(a'|a)$ are assumed to be strictly positive here, we can fix an arbitrary $a \in A$ and sum the above expression over all $a' \in A$ to obtain that $q/v(a)$ should be equal to 1. Since a was arbitrary, this means that the vector $v(a)$ is constant over a , and the result follows. \square

A.3 Proof of Corollary 2.3

First we note that, a careful examination of the proof of Theorem 2.1 part (i) shows that the entire argument remains valid for both cases $I_\pi(X_0; X_1) > 0$ and $I_\pi(X_0; X_1) = 0$, and also, without the positivity assumption on the transition matrix, as long as the chain \mathbf{X} is irreducible and aperiodic, since it is a standard exercise to show that the bivariate chain $\{Z_n\}$ is still ergodic in this case.

In view of the above remarks, we observe that the computation leading to (20) implies that, as $n \rightarrow \infty$,

$$\left(\sqrt{\frac{n}{\log \log n}} \right) D(\hat{P}_{X_0, n} \| \pi) = O \left(\sqrt{\frac{\log \log n}{n}} \right) = o(1), \quad \text{a.s.,}$$

and the same holds for each each of the three relative entropies in (19). Therefore, (21) becomes,

$$\begin{aligned} & \sqrt{\frac{n}{\log \log n}} [\hat{I}_n(P_{\bar{X}_0 \bar{X}_1}) - I_\pi(X_0; X_1)] \\ &= \frac{1}{\sqrt{n \log \log n}} \sum_{i=1}^n \left[\log \left(\frac{Q(X_i | X_{i-1})}{\pi(X_i)} \right) - I_\pi(X_0; X_1) \right] + o(1), \end{aligned}$$

a.s., as $n \rightarrow \infty$, and if $\sigma^2 > 0$ then an application of the law of the iterated logarithm [7] gives the claimed result. Finally, we note that, in view of the result in [26, Theorem 17.5.4], the same conclusion holds in the case $\sigma^2 = 0$. \square

A.4 Proof of Proposition 2.4

The first maximum in the definition of Δ_n , in view of (7), can be expressed as,

$$\max_{\theta \in \Theta} L_n(X_0^n; \theta) = \max_{\theta \in \Theta} \log \left(\prod_{i=1}^n Q_\theta(X_i | X_{i-1}) \right) = \max_Q \sum_{i=1}^n \log \left(Q(X_i | X_{i-1}) \right),$$

where the last maximum is over all transition matrices Q with all positive entries. Therefore,

$$\begin{aligned} \max_{\theta \in \Theta} L_n(X_0^n; \theta) &= \max_Q \sum_{a, a'} n \hat{P}_{X_0 X_1, n}(a, a') \log(Q(a' | a)) \\ &= -n \min_Q \left\{ D \left(\hat{P}_{X_0 X_1, n} \parallel Q \otimes \hat{P}_{X_0, n} \right) - \sum_{a, a'} \hat{P}_{X_0 X_1, n}(a, a') \log \left(\frac{\hat{P}_{X_0 X_1, n}(a, a')}{\hat{P}_{X_0, n}(a)} \right) \right\}, \end{aligned}$$

where $Q \otimes \hat{P}_{X, n}$ denotes the bivariate distribution $(Q \otimes \hat{P}_{X, n})(a, a') = \hat{P}_{X_0, n}(a) Q(a' | a)$. Clearly the above expression is minimized when the relative entropy term is zero, namely, when $Q(a' | a) = \hat{P}_{X_0 X_1, n}(a, a') / \hat{P}_{X_0, n}(a)$, so that,

$$\max_{\theta \in \Theta} L_n(X_0^n; \theta) = \sum_{a, a'} n \hat{P}_{X_0 X_1, n}(a, a') \log \left(\frac{\hat{P}_{X_0 X_1, n}(a, a')}{\hat{P}_{X_0, n}(a)} \right) = n[H(\hat{P}_{X_0, n}) - H(\hat{P}_{X_0 X_1, n})].$$

A similar (and simpler) computation, yields that the second maximum in (8) is,

$$\max_{\phi \in \Phi} L_n(X_0^n; h(\phi)) = -nH(\hat{P}_{X_0, n}).$$

Combining the last two equations with the definitions of Δ_n and $\hat{I}_n(P_{\bar{X}_0 \bar{X}_1})$ gives the claimed result. \square

A.5 Proof of Proposition 3.1

The existence of the entropy rate and the two expressions for $H(\mathbf{Y})$ in (i) can be established in several ways. For example, it is easy to check that the bivariate process $\{(X_n, Y_n)\}$ is asymptotically mean stationary (AMS) [13]. Then the univariate process $\{Y_n\}$ can be viewed as a

stationary coding of $\{(X_n, Y_n)\}$, and as such it is also AMS [12]. The results of (i) then follow immediately from the general results in [12].

For (ii), note that we can write,

$$\begin{aligned}
\sum_{i=1}^n H(Y_i | X_1^i, Y_1^{i-1}) &= \sum_{i=1}^k H(Y_i | X_1^i, Y_1^{i-1}) + \sum_{i=k+1}^n [H(X_i, Y_i | X_1^{i-1}, Y_1^{i-1}) - H(X_i | X_1^{i-1}, Y_1^{i-1})] \\
&= \sum_{i=1}^k H(Y_i | X_1^i, Y_1^{i-1}) + \sum_{i=k+1}^n [H(X_i, Y_i | X_{i-k}^{i-1}, Y_{i-k}^{i-1}) - H(X_i | X_{i-k}^{i-1}, Y_{i-k}^{i-1})] \\
&= \sum_{i=1}^k [H(Y_i | X_1^i, Y_1^{i-1}) - H(Y_i | X_{i-k}^i, Y_{i-k}^{i-1})] + \sum_{i=1}^n H(Y_i | X_{i-k}^i, Y_{i-k}^{i-1}). \quad (23)
\end{aligned}$$

By ergodicity and the continuity of the conditional entropy functional for finite-alphabet distributions, we have that $H(Y_i | X_{i-k}^i, Y_{i-k}^{i-1}) \rightarrow H(\bar{Y}_0 | \bar{X}_{-k}^0, \bar{Y}_{-k}^{-1})$ as $n \rightarrow \infty$. Therefore, dividing (23) by n and letting $n \rightarrow \infty$, the first term vanishes, and the Cesàro averages in the second term also converge,

$$\frac{1}{n} \sum_{i=1}^n H(Y_i | X_{i-k}^i, Y_{i-k}^{i-1}) \rightarrow H(\bar{Y}_0 | \bar{X}_{-k}^0, \bar{Y}_{-k}^{-1}).$$

The result in (ii) follows from this combined with (i) and the definition of the directed information rate.

Finally if, in addition, the process $\{Y_n\}$ is itself a Markov chain of order (no larger than) k , then its entropy rate is simply, $H(\bar{Y}_0 | \bar{Y}_{-k}^{-1})$, and the result in (iii) follows trivially from (ii). \square

A.6 Proof of Theorem 3.3

Both parts of the theorem will be established along the same lines as the proofs of the corresponding results in Theorem 2.1; for that reason, some minor details in the computations below will be omitted.

A.6.1 Proof of (i)

For the Gaussian convergence in part (i), recalling the definitions of the empirical $\hat{P}_{X_{-k}^0 Y_{-k}^0, n}$ and of the plug-in estimator in (9) and (10), respectively, we first express, $\hat{I}_n^{(k)}(\mathbf{X} \rightarrow \mathbf{Y})$ as a

mutual information,

$$\begin{aligned}
& \sum_{a_0^k, b_0^k} \hat{P}_{X_{-k}^0 Y_{-k}^0, n}(a_0^k, b_0^k) \log \left(\frac{\hat{P}_{X_{-k}^0 Y_0 | Y_{-k}^{-1}, n}(a_0^k, b_k | b_0^{k-1})}{\hat{P}_{X_{-k}^0 | Y_{-k}^{-1}, n}(a_0^k | b_0^{k-1}) \hat{P}_{Y_0 | Y_{-k}^{-1}, n}(b_k | b_0^{k-1})} \right) \\
&= \sum_{a_0^k, b_0^k} \hat{P}_{X_{-k}^0 Y_{-k}^0, n}(a_0^k, b_0^k) \log \left(\frac{\hat{P}_{X_{-k}^0 Y_{-k}^0, n}(a_0^k, b_0^k) \hat{P}_{Y_{-k}^{-1}, n}(b_0^{k-1})}{\hat{P}_{X_{-k}^0 Y_{-k}^{-1}, n}(a_0^k, b_0^{k-1}) \hat{P}_{Y_{-k}^0, n}(b_0^k)} \right) \\
&= \sum_{a_0^k, b_0^k} \hat{P}_{X_{-k}^0 Y_{-k}^0, n}(a_0^k, b_0^k) \log \left(\frac{P_{\bar{X}_{-k}^0 \bar{Y}_0 | \bar{Y}_{-k}^{-1}}(a_0^k, b_k | b_0^{k-1})}{P_{\bar{Y}_0 | \bar{Y}_{-k}^{-1}}(b_k | b_0^{k-1}) P_{\bar{X}_{-k}^0 | \bar{Y}_{-k}^{-1}}(a_0^k | b_0^{k-1})} \right) \tag{24} \\
&+ \sum_{a_0^k, b_0^k} \left\{ \hat{P}_{X_{-k}^0 Y_{-k}^0, n}(a_0^k, b_0^k) \right. \\
&\quad \left. \log \left(\frac{\hat{P}_{X_{-k}^0 Y_{-k}^0, n}(a_0^k, b_0^k) \hat{P}_{Y_0^{k-1}, n}(b_0^{k-1})}{\hat{P}_{X_{-k}^0 Y_{-k}^{-1}, n}(a_0^k, b_0^{k-1}) \hat{P}_{Y_{-k}^0, n}(b_0^k)} \cdot \frac{P_{\bar{X}_{-k}^0 \bar{Y}_{-k}^{-1}}(a_0^k, b_0^{k-1}) P_{\bar{Y}_{-k}^0}(b_0^k)}{P_{\bar{X}_{-k}^0 \bar{Y}_{-k}^0}(a_0^k, b_0^k) P_{\bar{Y}_{-k}^{-1}}(b_0^{k-1})} \right) \right\}. \tag{25}
\end{aligned}$$

Substituting the definition of $\hat{P}_{X_{-k}^0 Y_{-k}^0, n}$ in (24), simplifying, and expanding the logarithm in (25) a sum of four logarithms, we obtain that $\hat{I}_n^{(k)}(\mathbf{X} \rightarrow \mathbf{Y})$ equals,

$$\begin{aligned}
& \frac{1}{n} \sum_{i=1}^n \log \left(\frac{P_{\bar{X}_{-k}^0 \bar{Y}_0 | \bar{Y}_{-k}^{-1}}(X_{-k}^0, Y_0 | Y_{-k}^{-1})}{P_{\bar{Y}_0 | \bar{Y}_{-k}^{-1}}(Y_0 | Y_{-k}^{-1}) P_{\bar{X}_{-k}^0 | \bar{Y}_{-k}^{-1}}(X_{-k}^0 | Y_{-k}^{-1})} \right) + D \left(\hat{P}_{X_{-k}^0 Y_{-k}^0, n} \parallel P_{\bar{X}_{-k}^0 \bar{Y}_{-k}^0} \right) \\
&+ D \left(\hat{P}_{Y_{-k}^{-1}, n} \parallel P_{\bar{Y}_{-k}^{-1}} \right) - D \left(\hat{P}_{X_{-k}^0 Y_{-k}^{-1}, n} \parallel P_{\bar{X}_{-k}^0 \bar{Y}_{-k}^{-1}} \right) - D \left(\hat{P}_{Y_{-k}^0, n} \parallel P_{\bar{Y}_{-k}^0} \right). \tag{26}
\end{aligned}$$

We now claim that, each of the four relative entropies above, when multiplied by \sqrt{n} , converges to zero a.s. as $n \rightarrow \infty$. First recall that, as stated in the beginning of Section 3.2, the chain $\{Z_n = (X_{n-k}^n, Y_{n-k}^n)\}$ on $A^{k+1} \times B^{k+1}$ is ergodic, so we know that it satisfies the ergodic theorem, the central limit theorem, and the law of the iterated logarithm [7]. As we argued in the corresponding steps in the proof of Theorem 2.1, a quadratic Taylor expansion of the logarithm in the definition of $D \left(\hat{P}_{Y_{-k}^{-1}, n} \parallel P_{\bar{Y}_{-k}^{-1}} \right)$ gives,

$$\begin{aligned}
& \sqrt{n} D \left(\hat{P}_{Y_{-k}^{-1}, n} \parallel P_{\bar{Y}_{-k}^{-1}} \right) \\
&= -\sqrt{n} \sum_{b_0^{k-1} \in B^k} \hat{P}_{Y_{-k}^{-1}, n}(b_0^{k-1}) \left[\left(\frac{P_{\bar{Y}_{-k}^{-1}}(b_0^{k-1})}{\hat{P}_{Y_{-k}^{-1}, n}(b_0^{k-1})} - 1 \right) - \frac{1}{2} \left(\frac{P_{\bar{Y}_{-k}^{-1}}(b_0^{k-1})}{\hat{P}_{Y_{-k}^{-1}, n}(b_0^{k-1})} - 1 \right)^2 \frac{1}{\xi_n(b_0^{k-1})^2} \right] \\
&= \frac{\sqrt{n}}{2} \sum_{b_0^{k-1} \in B^k} \hat{P}_{Y_{-k}^{-1}, n}(b_0^{k-1}) \left(\frac{P_{\bar{Y}_{-k}^{-1}}(b_0^{k-1})}{\hat{P}_{Y_{-k}^{-1}, n}(b_0^{k-1})} - 1 \right)^2 \frac{1}{\xi_n(b_0^{k-1})^2} \\
&= \frac{\log \log n}{2\sqrt{n}} \sum_{b_0^{k-1} \in B^k} \left[\frac{1}{\hat{P}_{Y_{-k}^{-1}, n}(b_0^{k-1}) \xi_n(b_0^{k-1})^2} \left(\hat{P}_{Y_{-k}^{-1}, n}(b_0^{k-1}) - P_{\bar{Y}_{-k}^{-1}}(b_0^{k-1}) \right)^2 \left(\frac{n}{\log \log n} \right) \right],
\end{aligned}$$

for a (possibly random) $\xi_n(b_0^{k-1})$ between 1 and $P_{\bar{Y}_{-k}^{-1}}(b_0^{k-1})/\hat{P}_{Y_{-k}^{-1},n}(b_0^{k-1})$. Now, as in the proof of Theorem 2.1, the first term in the last sum above is bounded a.s. by the ergodic theorem, and the law of the iterated logarithm implies that the product of the next two terms,

$$\left(\hat{P}_{Y_{-k}^{-1},n}(b_0^{k-1}) - P_{\bar{Y}_{-k}^{-1}}(b_0^{k-1})\right)^2 \left(\frac{n}{\log \log n}\right) = \left\{ \frac{1}{\sqrt{n \log \log n}} \sum_{i=1}^n \left[\mathbb{I}_{\{Y_{i-k}^{i-1} = b_0^{k-1}\}} - P_{\bar{Y}_{-k}^{-1}}(b_0^{k-1}) \right] \right\}^2, \quad (27)$$

is also bounded a.s. So, each summand is a.s. bounded and, therefore, the entire expression tends to zero, a.s., as $n \rightarrow \infty$, as claimed.

Exactly the same argument shows that, after being multiplied by \sqrt{n} , the other three relative entropies in (26) also tend to zero, a.s., as $n \rightarrow \infty$, so that, a.s.,

$$\begin{aligned} & \sqrt{n} \left[\hat{I}_n^{(k)}(\mathbf{X} \rightarrow \mathbf{Y}) - I(\bar{Y}_0; \bar{X}_{-k}^0 | \bar{Y}_{-k}^{-1}) \right] \\ &= \frac{1}{\sqrt{n}} \sum_{i=1}^n \left[\log \left(\frac{P_{\bar{X}_{-k}^0 | \bar{Y}_0 | \bar{Y}_{-k}^{-1}}(X_{-k}^0, Y_0 | Y_{-k}^{-1})}{P_{\bar{Y}_0 | \bar{Y}_{-k}^{-1}}(Y_0 | Y_{-k}^{-1}) P_{\bar{X}_{-k}^0 | \bar{Y}_{-k}^{-1}}(X_{-k}^0 | Y_{-k}^{-1})} \right) - I(\bar{Y}_0; \bar{X}_{-k}^0 | \bar{Y}_{-k}^{-1}) \right] + o(1). \end{aligned} \quad (28)$$

The Gaussian convergence in part (i) follows by an application of the central limit theorem [7, Sec. I.16] to the above partial sums of a functional of the chain $\{Z_n\}$. The fact that the variance in (11) exists as the stated limit follows from [7, Theorem 3, p. 97], and the fact that it is finite is a direct consequence of the fact that the alphabets A and B are finite, and that the random variables $\{\log(\frac{\cdot}{\cdot})\}$ being summed in (28) are uniformly bounded. \square

A.6.2 Proof of (ii)

Recall the hypothesis testing setup of Section 3.3. In that notation, the class of all possible transition matrices $Q = Q_\theta$ is parametrized by an $m^k \ell^k (m\ell - 1)$ -dimensional vector,

$$\theta = (\theta_{i_1, i_2, \dots, i_k, j_1, j_2, \dots, j_k, i', j'}) = (\theta_{i_1^k, j_1^k, i', j'}),$$

where $1 \leq i_1, i_2, \dots, i_k \leq m$, $1 \leq j_1, j_2, \dots, j_k \leq \ell$, and $(i', j') \in A \times B$, $(i', j') \neq (m, \ell)$; we have used the same notation as before for strings of symbols, $i_1^k \in A^k$ and $j_1^k \in B^k$. In order for each θ to correspond to a transition probability matrix Q_θ with positive transitions, we restrict θ to the following open subset of $\mathbb{R}^{m^k \ell^k (m\ell - 1)}$,

$$\Theta = \left\{ \theta \in (0, 1)^{m^k \ell^k (m\ell - 1)} : \sum_{(i', j') \in A \times B, (i', j') \neq (m, \ell)} \theta_{i_1^k, j_1^k, i', j'} < 1, \text{ for all } i_1^k \in A^k, j_1^k \in B^k \right\},$$

so that the entries of the corresponding Q_θ are,

$$Q_\theta(i', j' | i_1^k, j_1^k) = \begin{cases} \theta_{i_1^k, j_1^k, i', j'}, & \text{if } (i', j') \neq (m, \ell), \\ 1 - \sum_{(u, v) \neq (m, \ell)} \theta_{i_1^k, j_1^k, u, v}, & \text{if } (i', j') = (m, \ell). \end{cases} \quad (29)$$

The null hypothesis is described by the parameter space $\Phi = \Gamma^x \times \Gamma^y$, where,

$$\begin{aligned}\Gamma^x &= \left\{ \gamma^x = (\gamma_{i_1^k, j_1^k, i'}^x) \in (0, 1)^{m^k \ell^k (m-1)} : \sum_{1 \leq i' \leq m-1} \gamma_{i_1^k, j_1^k, i'}^x < 1, \text{ for all } i_1^k \in A^k, j_1^k \in B^k \right\}, \\ \Gamma^y &= \left\{ \gamma^y = (\gamma_{j_1^k, j', j'}^y) \in (0, 1)^{\ell^k (\ell-1)} : \sum_{1 \leq j' \leq \ell-1} \gamma_{j_1^k, j', j'}^y < 1, \text{ for all } j_1^k \in B^k \right\}.\end{aligned}$$

Clearly Φ is an open set of $\mathbb{R}^{\ell^k [m^k (m-1) + (\ell-1)]}$ dimensions, which can be naturally embedded within Θ via the map $h : \Phi \rightarrow \Theta$, where each component of $h(\phi) = h(\gamma^x, \gamma^y)$ is,

$$h_{i_1^k, j_1^k, i', j'}(\phi) = \gamma_{i_1^k, j_1^k, i'}^x \cdot \gamma_{j_1^k, j', j'}^y, \quad (30)$$

for all components with $i' \neq m$ and $j' \neq \ell$, with the obvious extension to the two ‘edge’ cases (m, j') and (i', ℓ) .

In order to establish the χ^2 -convergence stated in the theorem, recall Proposition 3.5 which states that $2n$ times the plug-in estimator, $2n \hat{I}_n^{(k)}(\mathbf{X} \rightarrow \mathbf{Y})$, equals the log-likelihood ratio Δ_n defined in (18). The claimed result for the plug-in is an immediate consequence of the corresponding convergence result for Δ_n established by Billingsley in [5, Theorem 6.1], where $\Delta_n = \bar{\chi}_{t, t+1}^2(\hat{\phi})$ in the notation of [5].

To apply Billingsley’s result we only need to verify its main assumption, specifically Condition 6.1 on [5, p. 33], which requires that, throughout Φ , the matrix $Q_{h(\phi)}$ has continuous third order partial derivatives, and that the matrix $\mathcal{L}(\phi)$ defined below has rank $c = \ell^k [m^k (m-1) + (\ell-1)]$.

To that end, we note that, in the notation of [5], the size s of the alphabet $s = m\ell$, the memory length of the chain $t = k$, and the full parameter space $H_t = \Theta$ has dimension $r = s^{t+1} - s^t = m^k \ell^k (m\ell - 1)$. And since dimensionality c of the parameter space Φ for the null hypothesis is $c = \ell^k [m^k (m-1) + (\ell-1)]$, the limiting distribution of Δ_n has,

$$r - c = \ell^k (m^{k+1} - 1)(\ell - 1) \text{ degrees of freedom.}$$

By the definitions of Φ and the map h it is obvious that each component of the matrix $Q_{h(\phi)}$ has third-order partial derivatives with respect to every component of ϕ . Now consider the $s^{t+1} \times c = (m\ell)^{k+1} \times \ell^k [m^k (m-1) + (\ell-1)]$ matrix $\mathcal{L}(\phi)$, with entries given by the partial derivatives of each component $Q_{h(\phi)}(i', j' | i_1^k, j_1^k)$ of $Q_{h(\phi)}$ with respect to every component of $\phi = (\gamma^x, \gamma^y)$. Then the row of $\mathcal{L}(\phi)$ corresponding to (i_1^k, j_1^k, i', j') consists of all,

$$\frac{\partial Q_{h(\phi)}(i', j' | i_1^k, j_1^k)}{\partial \gamma_{u_1^k, v_1^k, u'}^x}, \quad (31)$$

followed by all,

$$\frac{\partial Q_{h(\phi)}(i', j' | i_1^k, j_1^k)}{\partial \gamma_{v_1^k, v'}^y}. \quad (32)$$

In view of (29) and (30), the derivatives in (31) are equal to zero unless $i_1^k = u_1^k$ and $j_1^k = v_1^k$, in which case,

$$\frac{\partial Q_{h(\phi)}(i', j' | i_1^k, j_1^k)}{\partial \gamma_{i_1^k, j_1^k, u'}^x} = \begin{cases} \gamma_{j_1^k, j'}^y, & \text{if } i' = u', i' \neq m \text{ and } j' \neq \ell, \\ -\gamma_{j_1^k, j'}^y, & \text{if } u' \neq m, i' = m \text{ and } j' \neq \ell, \\ 1 - \sum_{v \neq \ell} \gamma_{j_1^k, v}^y, & \text{if } i' = u', i' \neq m \text{ and } j' = \ell, \\ \sum_{v \neq \ell} \gamma_{j_1^k, v}^y - 1, & \text{if } u' \neq m, i' = m \text{ and } j' = \ell, \\ 0, & \text{in all other cases.} \end{cases}$$

Similarly, the derivatives in (32) are equal to zero unless $j_1^k = v_1^k$, in which case,

$$\frac{\partial Q_{h(\phi)}(i', j' | i_1^k, j_1^k)}{\partial \gamma_{v_1^k, v'}^y} = \begin{cases} \gamma_{i_1^k, j_1^k, i'}^x, & \text{if } j' = v', i' \neq m \text{ and } j' \neq \ell, \\ 1 - \sum_{u \neq m} \gamma_{i_1^k, j_1^k, u}^x, & \text{if } j' = v', i' = m \text{ and } j' \neq \ell, \\ -\gamma_{i_1^k, j_1^k, i'}^x, & \text{if } v' \neq \ell, i' \neq m \text{ and } j' = \ell, \\ \sum_{u \neq m} \gamma_{i_1^k, j_1^k, u}^x - 1, & \text{if } v' \neq \ell, i' = m \text{ and } j' = \ell, \\ 0, & \text{in all other cases.} \end{cases}$$

Since all the components of all the parameters γ^x and γ^y are strictly positive, a careful (and tedious) examination of the above expressions reveals that the matrix $\mathcal{L}(\phi)$ has rank $c = \ell^k[m^k(m-1) + (\ell-1)]$ throughout Φ . Therefore, Condition 6.1 of [5] holds as claimed, thus completing the proof. \square

A.7 Proof of Corollary 3.4

In view of Proposition 3.1, clearly it suffices to prove the first two assertions (13) and (14) of the corollary. We begin by noting that, as long as the pair process $\{(X_n, Y_n)\}$ is an irreducible and aperiodic Markov chain of order k , then a standard exercise to show that the chain $\{Z_n = (X_{n-k}^k, Y_{n-k}^k)\}$ is still ergodic. Then, an examination of the proof of Theorem 3.3 part (i) reveals that the entire argument remains valid for both cases $I(\bar{Y}_0; \bar{X}_{-k}^0 | \bar{Y}_{-k}^{-1}) > 0$ and $I(\bar{Y}_0; \bar{X}_{-k}^0 | \bar{Y}_{-k}^{-1}) = 0$, and without the positivity assumption on the transition matrix Q .

For the a.s.-rate in (13), we observe that the computation leading to (27) implies that, as $n \rightarrow \infty$,

$$\left(\sqrt{\frac{n}{\log \log n}} \right) D\left(\hat{P}_{Y_{-k}^{-1}, n}^{-1} \parallel P_{\bar{Y}_{-k}^{-1}} \right) = O\left(\sqrt{\frac{\log \log n}{n}} \right) = o(1), \quad \text{a.s.},$$

and that the same bound holds for each each of the four relative entropies in (26). Therefore, (28) becomes,

$$\begin{aligned} & \sqrt{\frac{n}{\log \log n}} \left[\hat{I}_n^{(k)}(\mathbf{X} \rightarrow \mathbf{Y}) - I(\bar{Y}_0; \bar{X}_{-k}^0 | \bar{Y}_{-k}^{-1}) \right] \\ &= \frac{1}{\sqrt{n \log \log n}} \sum_{i=1}^n \left[\log \left(\frac{P_{\bar{X}_{-k}^0 \bar{Y}_0 | \bar{Y}_{-k}^{-1}}(X_{-k}^0, Y_0 | Y_{-k}^{-1})}{P_{\bar{Y}_0 | \bar{Y}_{-k}^{-1}}(Y_0 | Y_{-k}^{-1}) P_{\bar{X}_{-k}^0 | \bar{Y}_{-k}^{-1}}(X_{-k}^0 | Y_{-k}^{-1})} \right) - I(\bar{Y}_0; \bar{X}_{-k}^0 | \bar{Y}_{-k}^{-1}) \right] \\ &+ o(1), \quad \text{a.s. as } n \rightarrow \infty. \end{aligned}$$

Then, if $\sigma^2 > 0$, the law of the iterated logarithm [7] implies (13) as claimed, and if $\sigma^2 = 0$, the same conclusion follows by [26, Theorem 17.5.4].

For the L^1 rate in (14), we first recall the expression for the plug-in estimator in (26) and claim that each of the four relative entropies there converge to zero at a rate $O(1/n)$ in L^1 . To see this, consider $D(\hat{P}_{Y_{-k}^{-1},n} \| P_{\bar{Y}_{-k}^{-1}})$; a first-order Taylor expansion for the logarithm in its definition gives,

$$D\left(\hat{P}_{Y_{-k}^{-1},n} \parallel P_{\bar{Y}_{-k}^{-1}}\right) = \sum_{b_0^{k-1}} \hat{P}_{Y_{-k}^{-1},n}(b_0^{k-1}) \left(\frac{\hat{P}_{Y_{-k}^{-1},n}(b_0^{k-1}) - P_{\bar{Y}_{-k}^{-1}}(b_0^{k-1})}{P_{\bar{Y}_{-k}^{-1}}(b_0^{k-1})} \right) \left(\frac{1}{\zeta_n(b_0^{k-1})} \right),$$

where, for those b_0^{k-1} for which $\hat{P}_{Y_{-k}^{-1},n}(b_0^{k-1})$ is nonzero, $\zeta_n(b_0^{k-1})$ is a (possibly random) constant between 1 and $\hat{P}_{Y_{-k}^{-1},n}(b_0^{k-1})/P_{\bar{Y}_{-k}^{-1}}(b_0^{k-1})$, while for the remaining b_0^{k-1} , $\zeta_n(b_0^{k-1})$ can be given an arbitrary (finite) value. Writing $S_n(b_0^{k-1})$ for the difference $\hat{P}_{Y_{-k}^{-1},n}(b_0^{k-1}) - P_{\bar{Y}_{-k}^{-1}}(b_0^{k-1})$, and $\rho_n(b_0^{k-1}) = [\zeta_n(b_0^{k-1}) - 1]P_{\bar{Y}_{-k}^{-1}}(b_0^{k-1})/S_n(b_0^{k-1})$, after some simple algebra we obtain,

$$D\left(\hat{P}_{Y_{-k}^{-1},n} \parallel P_{\bar{Y}_{-k}^{-1}}\right) = \sum_{b_0^{k-1}} S_n(b_0^{k-1}) \left(\frac{1 + S_n(b_0^{k-1})/P_{\bar{Y}_{-k}^{-1}}(b_0^{k-1})}{1 + \rho_n(b_0^{k-1})S_n(b_0^{k-1})/P_{\bar{Y}_{-k}^{-1}}(b_0^{k-1})} \right),$$

where now each $\rho_n(b_0^{k-1}) \in [0, 1]$. Using the simple inequality $(1+x)/(1+\rho x) \leq 1+x(1-\rho)$, which holds for all $x > -1$, for each $\rho \in [0, 1]$, gives,

$$D\left(\hat{P}_{Y_{-k}^{-1},n} \parallel P_{\bar{Y}_{-k}^{-1}}\right) \leq \sum_{b_0^{k-1}} S_n(b_0^{k-1}) \left(1 + [1 - \rho_n(b_0^{k-1})] \frac{S_n(b_0^{k-1})}{P_{\bar{Y}_{-k}^{-1}}(b_0^{k-1})} \right) \leq \sum_{b_0^{k-1}} \frac{S_n(b_0^{k-1})^2}{P_{\bar{Y}_{-k}^{-1}}(b_0^{k-1})},$$

because each ρ_n is between 0 and 1, and each S_n sums to zero by definition. In order to show that this relative entropy converges to zero in L^1 at a rate $O(1/n)$, it suffices to show that each term in the last sum does. And for that it suffices to show that for each b_0^{k-1} ,

$$nE[S_n(b_0^{k-1})^2] = O(1), \quad \text{as } n \rightarrow \infty. \quad (33)$$

But, as in (27), $S_n(b_0^{k-1})$ are simply the centered partial sums of a functional of the chain $\{Z_k\}$,

$$S_n(b_0^{k-1}) = \frac{1}{n} \sum_{i=1}^n \left[\mathbb{I}_{\{Y_{i-k}^{i-1} = b_0^{k-1}\}} - P_{\bar{Y}_{-k}^{-1}}(b_0^{k-1}) \right],$$

so the result of [7, Theorem 3, p. 97] tells us that actually, $nE[S_n(b_0^{k-1})^2]$ converges to a finite constant as $n \rightarrow \infty$. This implies (33), and establishes that,

$$E \left| D\left(\hat{P}_{Y_{-k}^{-1},n} \parallel P_{\bar{Y}_{-k}^{-1}}\right) \right| = O\left(\frac{1}{n}\right), \quad \text{as } n \rightarrow \infty.$$

The exact same argument shows that the same result also holds for the other three relative entropies in (26). Therefore, in order to establish the required result in (14) it suffices to show that,

$$E \left| \frac{1}{n} \sum_{i=1}^n \log \left(\frac{P_{\bar{X}_{-k}^0 \bar{Y}_0 | \bar{Y}_{-k}^{-1}}(X_{-k}^0, Y_0 | Y_{-k}^{-1})}{P_{\bar{Y}_0 | \bar{Y}_{-k}^{-1}}(Y_0 | Y_{-k}^{-1}) P_{\bar{X}_{-k}^0 | \bar{Y}_{-k}^{-1}}(X_{-k}^0 | Y_{-k}^{-1})} \right) - I(\bar{Y}_0; \bar{X}_{-k}^0 | \bar{Y}_{-k}^{-1}) \right| = O\left(\frac{1}{\sqrt{n}}\right). \quad (34)$$

But, once again, this can be seen as the L^1 norm of the centered partial sums of a functional of the chain $\{Z_n\}$. Therefore, Hölder's inequality combined with the result of [7, Theorem 3, p. 97], imply exactly (34), establishing (14) and completing the proof. \square

A.8 Proof of Proposition 3.5

We proceed along the same line as in the proof of Proposition 2.4. Recalling (17), the first maximum in the definition of Δ_n is,

$$\max_{\theta \in \Theta} L_n(X_{-k+1}^n, Y_{-k+1}^n; \theta) = \max_Q \sum_{i=1}^n \log \left(Q(X_i, Y_i | X_{i-k}^{i-1}, Y_{i-k}^{i-1}) \right),$$

where the last maximization is over all transition matrices Q with all positive entries, so that,

$$\begin{aligned} & \max_{\theta \in \Theta} L_n(X_{-k+1}^n, Y_{-k+1}^n; \theta) \\ &= \max_Q \sum_{a_0^k, b_0^k} n \hat{P}_{X_{-k}^0 Y_{-k}^0, n}(a_0^k, b_0^k) \log(Q(a_k, b_k | a_0^{k-1}, b_0^{k-1})) \\ &= -n \min_Q \left\{ D \left(\hat{P}_{X_{-k}^0 Y_{-k}^0, n}(a_0^k, b_0^k) \parallel Q \otimes \hat{P}_{X_{-k}^{-1} Y_{-k}^{-1}, n} \right) \right. \\ & \quad \left. - \sum_{a_0^k, b_0^k} \hat{P}_{X_{-k}^0 Y_{-k}^0, n}(a_0^k, b_0^k) \log \left(\frac{\hat{P}_{X_{-k}^0 Y_{-k}^0, n}(a_0^k, b_0^k)}{\hat{P}_{X_{-k}^{-1} Y_{-k}^{-1}, n}(a_0^{k-1}, b_0^{k-1})} \right) \right\}, \end{aligned}$$

where the distribution,

$$(Q \otimes \hat{P}_{X_{-k}^{-1} Y_{-k}^{-1}, n})(a_0^k, b_0^k) = \hat{P}_{X_{-k}^{-1} Y_{-k}^{-1}, n}(a_0^{k-1}, b_0^{k-1}) Q(a_k, b_k | a_0^{k-1}, b_0^{k-1}),$$

for all $a_0^k \in A^{k+1}$, $b_0^k \in B^{k+1}$. The above minimum is obviously achieved by making the relative entropy equal to zero, that is, by taking,

$$Q(a_k, b_k | a_0^{k-1}, b_0^{k-1}) = \hat{P}_{X_0 Y_0 | X_{-k}^{-1} Y_{-k}^{-1}, n}(a_k, b_k | a_0^{k-1}, b_0^{k-1}),$$

so that,

$$\max_{\theta \in \Theta} L_n(X_{-k+1}^n, Y_{-k+1}^n; \theta) = n \left[H \left(\hat{X}_{-k}^{-1}, \hat{Y}_{-k}^{-1} \right) - H \left(\hat{X}_{-k}^0, \hat{Y}_{-k}^0 \right) \right], \quad (35)$$

where, $(\hat{X}_{-k}^0, \hat{Y}_{-k}^0) \sim \hat{P}_{X_{-k}^0 Y_{-k}^0, n}$.

The computation for the second maximum in (18) is a little more involved, as it reduces to two different maximizations. But because both of these are very similar to the one just computed, we will give an outline of the steps involved without providing all the details. Recall the log-likelihood expression in (17) and that, under the null, Q admits the decomposition

in (16). We have:

$$\begin{aligned}
& \max_{\phi \in \Phi} L_n(X_{-k+1}^n, Y_{-k+1}^n; \theta) \\
&= \max_{Q^x, Q^y} \sum_{i=1}^n \log \left(Q^x(X_i | X_{i-k}^{i-1}, Y_{i-k}^{i-1}) Q^y(Y_i | Y_{i-k}^{i-1}) \right) \\
&= \max_{Q^x} \sum_{i=1}^n \log \left(Q^x(X_i | X_{i-k}^{i-1}, Y_{i-k}^{i-1}) \right) + \max_{Q^y} \sum_{i=1}^n \log \left(Q^y(Y_i | Y_{i-k}^{i-1}) \right) \\
&= \max_{Q^x} \sum_{a_0^k, b_0^{k-1}} n \hat{P}_{X_{-k}^0 Y_{-k}^{-1}, n}(a_0^k, b_0^{k-1}) \log(Q^x(a_k | a_0^{k-1}, b_0^{k-1})) \\
&\quad + \max_{Q^y} \sum_{b_0^k} n \hat{P}_{Y_{-k}^0, n}(b_0^k) \log(Q^y(b_k | b_0^{k-1})) \\
&= n \sum_{a_0^k, b_0^{k-1}} \hat{P}_{X_{-k}^0 Y_{-k}^{-1}, n}(a_0^k, b_0^{k-1}) \log \left(\frac{\hat{P}_{X_{-k}^0 Y_{-k}^{-1}, n}(a_0^k, b_0^{k-1})}{\hat{P}_{X_{-k}^{-1} Y_{-k}^{-1}, n}(a_0^{k-1}, b_0^{k-1})} \right) \\
&\quad + n \sum_{b_0^{k-1}} \hat{P}_{Y_{-k}^0, n}(b_0^k) \log \left(\frac{\hat{P}_{Y_{-k}^0, n}(b_0^k)}{\hat{P}_{Y_{-k}^{-1}, n}(b_0^{k-1})} \right) \\
&= n \left[-H(\hat{X}_{-k}^0, \hat{Y}_{-k}^{-1}) + H(\hat{X}_{-k}^{-1}, \hat{Y}_{-k}^{-1}) - H(\hat{Y}_{-k}^0) + H(\hat{Y}_{-k}^{-1}) \right]. \tag{36}
\end{aligned}$$

Combining (35) and (36) and using the chain rule,

$$\begin{aligned}
\Delta_n &= 2 \left\{ \max_{\theta \in \Theta} L_n(X_{-k+1}^n, Y_{-k+1}^n; \theta) - \max_{\phi \in \Phi} L_n(X_{-k+1}^n, Y_{-k+1}^n; h(\phi)) \right\} \\
&= 2n \left[H(\hat{Y}_0 | \hat{Y}_{-k}^{-1}) - H(\hat{Y}_0 | \hat{X}_{-k}^0, \hat{Y}_{-k}^{-1}) \right],
\end{aligned}$$

which, recalling the definition of $\hat{I}_n^{(k)}(\mathbf{X} \rightarrow \mathbf{Y})$, is precisely the claimed result. \square

Acknowledgments

We wish to thank Wojtek Szpankowski and Vasiliki Gioutlaki for interesting discussions during the preliminary phase of this work. We are also grateful to the three anonymous reviewers for their very detailed, insightful and useful comments.

References

- [1] H.M. Aktulga, I. Kontoyiannis, L.A. Lyznik, L. Szpankowski, A.Y. Grama, and W. Szpankowski. Identifying statistical dependence in genomic sequences via mutual information estimates. *EURASIP J. Bioinformatics Syst. Biol.*, 2007:3:1–3:11, July 2007.
- [2] P-O. Amblard and O.J.J. Michel. The relation between Granger causality and directed information theory: A review. *Entropy*, 15(1):113, 2012.
- [3] R.P. Anderson and M. Porfiri. Assessing significance of information flow in high dimensional dynamical systems with few data. In *ASME 2014 Dynamic Systems and Control Conference*, pages V002T24A002–V002T24A002. American Society of Mechanical Engineers, 2014.
- [4] L. Barnett and T. Bossomaier. Transfer entropy as a log-likelihood ratio. *Physical review letters*, 109(13):138105, 2012.
- [5] P. Billingsley. *Statistical Inference for Markov Processes*. Statistical Research Monographs. University of Chicago Press, 1961.
- [6] L. Le Cam. Asymptotic methods in statistical decision theory. 1986.
- [7] K.L. Chung. *Markov Chains with Stationary Transition Probabilities*. Springer-Verlag, New York, 1967.
- [8] R. Dabora and A.J. Goldsmith. The capacity region of the degraded finite-state broadcast channel. *Information Theory, IEEE Transactions on*, 56(4):1828–1851, April 2010.
- [9] B. Goebel, Z. Dawy, J. Hagenauer, and J.C. Mueller. An approximation to the distribution of finite sample size mutual information estimates. In *Proc. of the IEEE International Conference on Communications (ICC 2005)*, Seoul, South Korea, May 2005.
- [10] C. Gourieroux, A. Monfort, and E. Renault. Kullback causality measures. *Annales d'Économie et de Statistique*, no. 6/7:369–410, 1987.
- [11] C.W.J. Granger. Investigating causal relations by econometric models and cross-spectral methods. *Econometrica*, 37(3):424–438, 1969.
- [12] R.M. Gray. *Entropy and Information Theory*. Springer-Verlag, New York, 1990.
- [13] R.M. Gray and J.C. Kieffer. Asymptotically mean stationary measures. *Ann. Probab.*, 8(5):962–973, 1980.
- [14] D. Gunduz and E. Erkip. Lossless transmission of correlated sources over a multiple access channel with side information. In *Data Compression Conference, 2007. DCC '07*, pages 83–92, March 2007.

- [15] J. Hagenauer and J.C. Mueller. Genomic analysis using methods from information theory. In *Proc. of the Inform. Theory Workshop*, pages 55–59, San Antonio, TX, October 2004.
- [16] Y. Han, J. Jiao, and T. Weissman. Minimax estimation of discrete distributions. In *2015 IEEE International Symposium on Information Theory (ISIT) Proceedings*, pages 2291–2295. IEEE, 2015.
- [17] J. Haskovec. A note on the consensus finding problem in communication networks with switching topologies. *Applicable Analysis*, 94(5):991–998, 2015.
- [18] J. Jiao, H.H. Permuter, L. Zhao, Y.-H. Kim, and T. Weissman. Universal estimation of directed information. *Information Theory, IEEE Transactions on*, 59(10):6220–6242, Oct 2013.
- [19] Y-H. Kim. A coding theorem for a class of stationary channels with feedback. *Information Theory, IEEE Transactions on*, 54(4):1488–1499, April 2008.
- [20] I. Kontoyiannis. Asymptotic recurrence and waiting times for stationary processes. *J. Theoret. Probab.*, 11:795–811, 1998.
- [21] G. Kramer. *Directed Information for Channels With Feedback*. PhD thesis, Swiss Fed. Inst. Technol. (ETH), Zurich, Switzerland, 1998.
- [22] G. Kramer. Information networks with in-block memory. *Preprint, arXiv:1206.5389*, 2014.
- [23] E.L. Lehmann and J.P. Romano. *Testing Statistical Hypotheses*. Springer Texts in Statistics. Springer, 2005.
- [24] H Marko. The bidirectional communication theory: A generalization of information theory. *IEEE Trans. Comm.*, 21(1):1335–1351, 1973.
- [25] J. Massey. Rate-distortion in near-linear time. In *International Symposium on Information Theory and its Applications*, pages 303–305, November 1990.
- [26] S.P. Meyn and R.L. Tweedie. *Markov Chains and Stochastic Stability*. Springer-Verlag, London, 1993.
- [27] E.G. Miller. A new class of entropy estimators for multi-dimensional densities. In *Acoustics, Speech, and Signal Processing, 2003. Proceedings. (ICASSP '03). 2003 IEEE International Conference on*, volume 3, pages III–297–300 vol.3, April 2003.
- [28] Y. Ogata and N. Inagaki. The weak convergence of the likelihood ratio random fields for Markov observations. *Annals of the Institute of Statistical Mathematics*, 29(1):165–187, 1977.
- [29] H.H. Permuter, Y-H. Kim, and T. Weissman. Interpretations of directed information in portfolio theory, data compression, and hypothesis testing. *Information Theory, IEEE Transactions on*, 57(6):3248–3259, June 2011.
- [30] H.H. Permuter, T. Weissman, and A.J. Goldsmith. Finite state channels with time-invariant deterministic feedback. *Information Theory, IEEE Transactions on*, 55(2):644–662, Feb 2009.

- [31] C.J. Quinn, T.P. Coleman, N. Kiyavash, and N.G. Hatsopoulos. Estimating the directed information to infer causal relationships in ensemble neural spike train recordings. *Journal of Computational Neuroscience*, 30(1):17–44, 2011.
- [32] C.J. Quinn, J. Etesami, N. Kiyavash, and T.P. Coleman. Robust directed tree approximations for networks of stochastic processes. In *2013 IEEE International Symposium on Information Theory (ISIT) Proceedings*, pages 2254–2258. IEEE, 2013.
- [33] M.A. Rahimian and V.M. Preciado. Detection and isolation of link failures under the agreement protocol. In *Decision and Control (CDC), 2013 IEEE 52nd Annual Conference on*, pages 7364–7369, Dec 2013.
- [34] A. Rao, A.O. Hero, D.J. States, and J.D. Engel. Using directed information to build biologically relevant influence networks. *Journal of Bioinformatics and Computational Biology*, 06(03):493–519, 2008.
- [35] M.J. Schervish. *Theory of Statistics*. Springer Series in Statistics. Springer, New York, 1996.
- [36] B. Shrader and H. Permuter. Feedback capacity of the compound channel. *Information Theory, IEEE Transactions on*, 55(8):3629–3644, Aug 2009.
- [37] E.I. Silva, M.S. Derpich, J. Ostergaard, and M.A. Encina. A characterization of the minimal average data rate that guarantees a given closed-loop performance level. *Preprint, available at <http://arxiv.org/abs/1407.0118>*, 2013.
- [38] S.C. Tatikonda. *Control Under Communication Constraints*. PhD thesis, Massachusetts Inst. Technol., Cambridge, MA, 2000.
- [39] A.W. van der Vaart. *Asymptotic Statistics*. Cambridge Series in Statistical and Probabilistic Mathematics. Cambridge University Press, 2000.
- [40] R. Venkataramanan and S.S. Pradhan. Source coding with feed-forward: Rate-distortion theorems and error exponents for a general source. *Information Theory, IEEE Transactions on*, 53(6):2154–2179, June 2007.
- [41] T. Weissman, Y-H. Kim, and H.H. Permuter. Directed information, causal estimation, and communication in continuous time. *Information Theory, IEEE Transactions on*, 59(3):1271–1287, March 2013.
- [42] S.S. Wilks. The large-sample distribution of the likelihood ratio for testing composite hypotheses. *Ann. Math. Statist.*, 9(1):60–62, March 1938.
- [43] F.M.J. Willems, Y.M. Shtarkov, and T.J. Tjalkens. Context tree weighting: Basic properties. *IEEE Trans. Inform. Theory*, 41(3):653–664, 1995.
- [44] A.A. Yushkevich. On limit theorems connected with the concept of the entropy of Markov chains. *Uspehi Mat. Nauk*, 8:177–180, 1953. (Russian).
- [45] A.A. Zaidi, T.J. Oechtering, S. Yüksel, and M. Skoglund. Stabilization of linear systems over gaussian networks. *Automatic Control, IEEE Transactions on*, 59(9):2369–2384, Sept 2014.